

THE HOLOMORPHIC COUCH THEOREM

MAXIME FORTIER BOURQUE

ABSTRACT. We prove that if two conformal embeddings between Riemann surfaces with finite topology are homotopic, then they are isotopic through conformal embeddings. Furthermore, we show that the space of all conformal embeddings in a given homotopy class deformation retracts into a point, a circle, a torus, or the unit tangent bundle of the codomain, depending on the induced homomorphism on fundamental groups. Quadratic differentials play a central role in the proof.

1. INTRODUCTION

Loosely speaking, the 1-*parametric h-principle* is said to hold for a class of maps between manifolds if the only obstructions to connecting two maps in the class through maps in the same class are topological [10, p.60]. For example, the 1-parametric *h-principle* holds for immersions of S^2 in \mathbb{R}^3 , so that the standard sphere can be turned inside out via immersions. This is known as Smale’s paradox. Of course, there are situations where the 1-parametric *h-principle* fails due to geometric obstructions. A famous example is Gromov’s symplectic camel theorem, which says that one cannot move a closed 4-dimensional ball through a hole in a wall in \mathbb{R}^4 via symplectic embeddings if the ball is bigger than the hole.

In this paper, we prove that the 1-parametric *h-principle* holds for conformal embeddings between finite Riemann surfaces, where a *finite Riemann surface* is a Riemann surface with finitely generated fundamental group.

Theorem 1.1 (The holomorphic couch theorem). *If two conformal embeddings between finite Riemann surfaces are homotopic, then they are isotopic through conformal embeddings.*

We think of the codomain as a house and the domain as a couch that we want to move around in the house without changing its holomorphic structure. Hence the name “holomorphic couch”.

Given finite Riemann surfaces X and Y , and a topological embedding $h : X \rightarrow Y$, we define $\text{CEmb}(X, Y, h)$ to be the set of all conformal embeddings homotopic to h . We equip this set with the compact-open topology,

2010 *Mathematics Subject Classification.* 30F60, 30C75, 32H02.

Key words and phrases. conformal, quasiconformal, embedding, homotopy, isotopy, *h-principle*, Teichmüller’s theorem, quadratic differential, extremal length.

Research supported by the Natural Sciences and Engineering Research Council of Canada and NSF grant DMS 0905812.

which is the same as the topology of uniform convergence on compact sets with respect to any metric inducing the correct topology on Y . Theorem 1.1 is equivalent¹ to the statement that $\text{CEmb}(X, Y, h)$ is path-connected whenever it is non-empty.

In fact, we prove a stronger result. Namely, we determine the homotopy type of the space $\text{CEmb}(X, Y, h)$. The answer depends on the image of h at the level of fundamental groups. We say that h is *trivial*, *cyclic*, or *generic* if the image of the induced homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ is trivial, infinite cyclic, or non-abelian, respectively.

Theorem 1.2 (Strong holomorphic couch). *Suppose that $\text{CEmb}(X, Y, h)$ is non-empty. Then $\text{CEmb}(X, Y, h)$ is homotopy equivalent to either the unit tangent bundle of Y , a circle, or a point, depending on whether h is trivial, cyclic, or generic. This is unless Y is a torus and h is non-trivial, in which case $\text{CEmb}(X, Y, h)$ is homotopy equivalent to a torus.*

If h is generic, then $\text{CEmb}(X, Y, h)$ is contractible. The rest of the introduction gives an outline of the proof for this case.

Ioffe's theorem. A *half-translation structure* on a finite Riemann surface is an atlas of conformal charts (for the surface minus a finite set) whose transition functions are of the form $z \mapsto \pm z + c$ (with some extra condition on the finite set) and in which the surface has finite area. This data can be specified by an integrable holomorphic quadratic differential. The charts of the half-translation structure should also map the ideal boundary to curves which are piecewise horizontal or vertical.

A *Teichmüller embedding* is an injective map $f : X \rightarrow Y$ for which there exist half-translation structures on X and Y , and a constant $K \geq 1$ such that:

- f is locally of the form $x + iy \mapsto Kx + iy$ (up to sign and translation);
- the ideal boundary of X and Y is horizontal;
- $Y \setminus f(X)$ is a finite union of points and horizontal arcs.

Our main tool is a theorem of Ioffe which says that extremal quasiconformal embeddings and Teichmüller embeddings are one and the same [20].

Theorem 1.3 (Ioffe). *Let $f : X \rightarrow Y$ be a quasiconformal embedding which is not conformal. Then f has minimal dilatation (i.e. is closest to being conformal) in its homotopy class if and only if it is a Teichmüller embedding.*

In the special case where X and Y are closed, we recover Teichmüller's celebrated theorem, since an embedding between closed surfaces is a homeomorphism. What is different from Teichmüller's theorem, however, is that Teichmüller embeddings are not necessarily unique in their homotopy class. This is an important issue which we discuss next.

¹Since X is locally compact Hausdorff, a path $[0, 1] \rightarrow \text{Map}(X, Y)$ is the same as a homotopy $X \times [0, 1] \rightarrow Y$ [29, p.287].

Slit mappings. A *slit mapping* is a conformal Teichmüller embedding, i.e. one with $K = 1$. In this case, the half-translation structure on X is redundant, for it is the pull-back of the half-translation structure on Y by the slit mapping.

We show that if the space $\text{CEmb}(X, Y, h)$ contains a slit mapping, then every element of $\text{CEmb}(X, Y, h)$ is a slit mapping and $\text{CEmb}(X, Y, h)$ is naturally homeomorphic to a point or a compact interval.

Theorem 1.4 (Slit mappings are almost rigid). *Suppose that $\text{CEmb}(X, Y, h)$ contains a slit mapping with respect to a half-translation structure Ψ on Y . Then every $g \in \text{CEmb}(X, Y, h)$ is a slit mapping with respect to Ψ . Moreover, for every $x \in X$, the evaluation map $\text{CEmb}(X, Y, h) \rightarrow Y$ sending g to $g(x)$ is a homeomorphism onto its image, which is a (possibly degenerate) compact horizontal arc whose length does not depend on x .*

A better way to say this is: any conformal embedding homotopic to a slit mapping differs from the latter by a horizontal translation which can be performed gradually.

Since an interval is contractible, Theorem 1.4 implies Theorem 1.2 whenever the space $\text{CEmb}(X, Y, h)$ contains a slit mapping. The idea is then to reduce the general case to the above one by enlarging X until it barely fits in Y .

Modulus of extension. Suppose that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Then X has at least one ideal boundary component². Thus, we may define a 1-parameter family of enlargements of X as follows. We first choose an analytic parametrization $S^1 \rightarrow C$ of each ideal boundary component C of X . Then, for every $r \in (0, \infty]$, we let X_r be the bordered surface $X \cup \partial X$ with a copy of the cylinder $S^1 \times [0, r)$ glued to each ideal boundary component along $S^1 \times \{0\}$ via the fixed parametrization. We also let $X_0 = X$.

For any $f \in \text{CEmb}(X, Y, h)$, we define the *modulus of extension* $\mathfrak{m}(f)$ as the supremum of the set of $r \in [0, \infty]$ such that f extends to a conformal embedding of X_r into Y . Montel's theorem in complex analysis implies that:

- for every $f \in \text{CEmb}(X, Y, h)$ the supremum $\mathfrak{m}(f)$ is achieved by a unique conformal embedding $\hat{f} : X_{\mathfrak{m}(f)} \rightarrow Y$ extending f ;
- $\text{CEmb}(X, Y, h)$ is compact;
- \mathfrak{m} is upper semi-continuous.

In particular, \mathfrak{m} achieves its maximum value. Using Ioffe's theorem, it is not too hard to show that if \mathfrak{m} attains its maximum at f , then \hat{f} is a slit mapping. We prove that the same holds if f is a local maximum of \mathfrak{m} .

²If not, then for any $f \in \text{CEmb}(X, Y, h)$ the complement $Y \setminus f(X)$ is finite so that f is a slit mapping, as long as Y supports a half-translation structure. The only finite Riemann surface with non-abelian fundamental group which does not support any half-translation structure is the triply punctured sphere. But if each of X and Y is a punctured sphere and h is generic, then $\text{CEmb}(X, Y, h)$ has at most one element, so there is nothing to show.

Theorem 1.5 (Characterization of local maxima). *Let f be a local maximum of \mathfrak{m} , and let \hat{f} be the maximal conformal extension of f to $X_{\mathfrak{m}(f)}$. Then \hat{f} is a slit mapping. Conversely, if $g : X_r \rightarrow Y$ is a slit mapping, then $g|_X$ is a global maximum of \mathfrak{m} . The set M of all local maxima of \mathfrak{m} is homeomorphic to a point or a compact interval, and \mathfrak{m} is constant on M .*

The initial motivation for studying \mathfrak{m} was to think of it as a Morse function for the space $\text{CEmb}(X, Y, h)$. In an ideal world, flowing along the gradient of \mathfrak{m} would yield a deformation retraction of $\text{CEmb}(X, Y, h)$ into M . This does not make sense though, since \mathfrak{m} is not even continuous unless it is constant equal to zero. However, the connectedness of $\text{CEmb}(X, Y, h)$ is an easy consequence of Theorem 1.5.

Theorem 1.6 (Weak holomorphic couch). *$\text{CEmb}(X, Y, h)$ is connected.*

To improve upon this, we show that there are no local obstructions to contractibility.

Where can one point go? Given a point $x \in X$, we are interested in set of points in Y where x can be mapped by the elements of $\text{CEmb}(X, Y, h)$. It is convenient to also keep track of how x gets mapped to a given $y \in Y$ in the following sense. If $f \in \text{CEmb}(X, Y, h)$, then by definition there exists a homotopy

$$H : X \times [0, 1] \rightarrow Y$$

from h to f . The homotopy class of the path $t \mapsto H(x, t)$ from $h(x)$ to $f(x)$ does not depend on the particular choice of H . We thus let $\text{lift}_x(f)$ be the homotopy class of this path. This is an element of the universal cover of Y , which can be identified with the unit disk \mathbb{D} .

The map $\text{lift}_x : \text{CEmb}(X, Y, h) \rightarrow \mathbb{D}$ is continuous, and we let $\text{Blob}(x)$ be its image. The blob is simpler than the image of the evaluation map in much the same way as Teichmüller space is simpler than moduli space. Indeed, the blob is as simple as can be.

Theorem 1.7 (The blob is a disk). *If $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping, then $\text{Blob}(x)$ is homeomorphic to a closed disk.*

The proof has four steps:

- (1) the blob is compact and connected;
- (2) the blob is semi-smooth;
- (3) every semi-smooth set is a 2-manifold with boundary;
- (4) there are no holes in the blob.

Step (1) follows because $\text{Blob}(x)$ is the continuous image of $\text{CEmb}(X, Y, h)$ which is compact and connected. The next two steps require some terminology. We say that a closed subset B of \mathbb{R}^2 is *semi-smooth* if:

- at each boundary point of B , there is at least one *normal vector*, i.e. a non-zero vector pointing orthogonally away from B ;

- the set of all normal vectors at a point is convex (and is thus a cone with total angle less than π);
- any non-zero limit of normal vectors (not necessarily based at the same point) is a normal vector.

For example, 2-manifolds with C^1 boundary are semi-smooth, and so are closed convex sets with non-empty interior. To be clear, semi-smooth sets are allowed to have corners, but these should look convex as we zoom in on them. Corners pointing inwards—or dents—are forbidden since the first condition is not met. Similarly, outward pointing cusps are forbidden by the second condition.

We now describe some of the ideas involved in step (2). Let $\pi : \mathbb{D} \rightarrow Y$ be the universal covering map, let $\tilde{y} = \text{lift}_x(f)$ be on the boundary of $\text{Blob}(x)$, and let $y = \pi(\tilde{y}) = f(x)$. It follows from Ioffe's theorem and a limiting argument that the restriction $f : X \setminus x \rightarrow Y \setminus y$ is a slit mapping. We say that f is a *slit mapping rel x* . We use the notion of extremal length, and in particular Gardiner's formula, to determine which tangent vectors are normal to $\text{Blob}(x)$ at \tilde{y} . If \mathcal{F} denotes the horizontal foliation of a half-translation structure on $X \setminus x$ with respect to which f is a slit mapping rel x , then the extremal length of \mathcal{F} on $X \setminus x$ is equal to the extremal length of $f_*\mathcal{F}$ on $Y \setminus y$. Now the point is that extremal length does not increase under conformal embeddings. Therefore, whenever a point y' close enough to y is such that the extremal length of $f_*\mathcal{F}$ is bigger on $Y \setminus y'$ than on $Y \setminus y$, then the lift of y' nearest to \tilde{y} is outside $\text{Blob}(x)$. But Gardiner's formula tells us precisely how the extremal length of $f_*\mathcal{F}$ on $Y \setminus y'$ varies as we move y' around. The end result is that \mathbf{v} is normal to $\text{Blob}(x)$ at \tilde{y} if and only if $\pi_*(\mathbf{v})$ is vertical for some half-translation structure on $Y \setminus y$ with respect to which f is a slit mapping rel x . We deduce that the blob is semi-smooth.

It is well-known that every closed convex set with non-empty interior in \mathbb{R}^2 is a 2-manifold with boundary. Step (3) is a generalization of this to semi-smooth sets.

Given (1), (2) and (3), we have that $\text{Blob}(x)$ is a connected compact 2-manifold with boundary in the plane. It only remains to show that $\text{Blob}(x)$ does not have any holes. For this, we consider the family of blobs corresponding to the enlargements of X constructed earlier. In other words, we look at $\text{Blob}(x, X_r)$ for $r \in [0, R]$, where R is the maximal value of the modulus of extension \mathbf{m} . These blobs are nested and are all 2-manifolds with boundary except for $\text{Blob}(x, X_R)$, which is either a point or a compact interval. We show that the number of connected components of $\mathbb{D} \setminus \text{Blob}(x, X_r)$ does not depend on r . Since the complement of $\text{Blob}(x, X_R)$ is connected, it follows that $\mathbb{D} \setminus \text{Blob}(x)$ is connected and hence that $\text{Blob}(x)$ is a disk.

Moving one point at a time. Pick a countable dense set $\{x_1, x_2, \dots\}$ in X and let $F \in \text{CEmb}(X, Y, h)$ be any conformal embedding. We now explain how to get a deformation retract of $\text{CEmb}(X, Y, h)$ into $\{F\}$ by moving one point at a time. Given $f \in \text{CEmb}(X, Y, h)$, we join $\text{lift}_{x_1}(f)$ to $\text{lift}_{x_1}(F)$ by

a path γ_1 in $\text{Blob}(x_1)$. Such a path exists since $\text{Blob}(x_1)$ is homeomorphic to a closed disk.

For every $t \in [0, 1]$, we then look at where x_2 can go under maps $g \in \text{CEmb}(X, Y, h)$ which satisfy $\text{lift}_{x_1}(g) = \gamma_1(t)$. This defines a new kind of blob, call it $\text{Blob}_t(x_2)$. We find that $\text{Blob}_t(x_2)$ moves continuously with t , which allows us to construct a second path γ_2 from $\text{lift}_{x_2}(f)$ to $\text{lift}_{x_2}(F)$ with the property that for every $t \in [0, 1]$, the point $\gamma_2(t)$ belongs to $\text{Blob}_t(x_2)$.

Proceeding by induction, we obtain a sequence of paths $\gamma_1, \gamma_2, \dots$ such that for every $n \in \mathbb{N}$ and every $t \in [0, 1]$, there exists at least one map $f_t^n \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_{x_j}(f_t^n) = \gamma_j(t)$ for every $j \in \{1, \dots, n\}$. If we fix t and pass to a subsequence, we get some limit $f_t \in \text{CEmb}(X, Y, h)$ for which $\text{lift}_{x_j}(f_t) = \gamma_j(t)$ for every $j \in \mathbb{N}$. Since any two limits agree on the dense set $\{x_1, x_2, \dots\}$, we actually have convergence $f_t^n \rightarrow f_t$ for the whole sequence. By a similar argument, f_t depends continuously on t . We thus found a path from f to F in $\text{CEmb}(X, Y, h)$.

We construct the paths $\gamma_1, \gamma_2, \dots$ carefully enough so that they depend continuously on the initial map f . The end result is a deformation retraction of $\text{CEmb}(X, Y, h)$ into $\{F\}$.

Notes and references. The holomorphic couch problem arose in the context of renormalization in complex dynamics. Although our theorem does not have any direct applications to dynamics, some of the tools used here—such as extremal length and Ioffe’s theorem—do (see [39]).

The space $\text{CEmb}(\mathbb{D}, \mathbb{C})$, or rather its subspace \mathcal{S} of conformal embeddings $f : \mathbb{D} \rightarrow \mathbb{C}$ satisfying $f(0) = 0$ and $f'(0) = 1$, was the subject of much interest until the solution of the Bieberbach conjecture by de Branges in 1984. It is easy to see that \mathcal{S} is contractible. On the other hand, \mathcal{S} has isolated points when equipped with the topology of uniform convergence of the Schwarzian derivative instead of the compact-open topology [40]. The literature on the class \mathcal{S} is quite vast. In comparison, not much has been written about conformal embeddings between general Riemann surfaces. Exceptions include [34], [20], [21], [6], and [35].

The holomorphic couch problem for embeddings of a multiply punctured disk in a multiply punctured sphere was first considered in [32]. However, the solution presented there relies in part on a rigidity claim [31] which is known to be false in general [22]. On a related note, Ioffe asserts that generic slit mappings are completely rigid in [20], but this is wrong. Indeed, we give examples where $\text{CEmb}(X, Y, h)$ is an interval and not a singleton in Section 3. The conclusion of Theorem 1.4 is thus optimal.

There are various functionals similar to the modulus of extension \mathfrak{m} for which analogous versions of Theorem 1.5 hold. For example, instead of gluing the annuli to X , one can keep them disjoint from X [6]. More generally, there is a plethora of extremal problems on Riemann surfaces whose solutions involve quadratic differentials (see e.g. [23] and the references therein).

For the class \mathcal{S} of normalized univalent functions from \mathbb{D} to \mathbb{C} , a suitable version of the blob is actually a round disk. More precisely, for every $z \in \mathbb{D}$ and every $f \in \mathcal{S}$, the quantity $w = \log(f(z)/z)$ satisfies

$$\left| w - \log \frac{1}{1 - |z|^2} \right| \leq \log \frac{1 + |z|}{1 - |z|}$$

and every value w satisfying the inequality is achieved for some $f \in \mathcal{S}$. This was proved by Grunsky in 1932 (see [5, p.323]). The blob for K -quasiconformal homeomorphisms was studied in [38] and [9]. Our approach for proving that the blob is homeomorphic to a closed disk seems similar to Strebel's, although the context is different.

The idea of moving one point at a time to get an isotopy is reminiscent of the finite "holomorphic axiom of choice" used by Ślodkowski to extend holomorphic motions [36]. Our isotopies are holomorphic in the space variable and continuous in the time variable, whereas holomorphic motions are the other way around.

In closing, we note that our results can be extended to the situation where the domain X is a finite union of finite Riemann surfaces. This is the setup considered in Ioffe's paper [20]. We restrict ourselves to connected domains for simplicity.

Acknowledgements. This is my Ph.D. thesis completed at the Graduate Center of the City University of New York. I thank my advisor Jeremy Kahn for suggesting this problem and for his continued interest and support. I also thank Dylan Thurston, Kevin Pilgrim, Frederick Gardiner, Patrick Hooper, Nick Spizzirri, and Joe Adams for many helpful conversations related to this work, François Labourie for pointing out the connection with the h -principle, and Paul Carter for his careful reading of part of the manuscript.

CONTENTS

1. Introduction	1
2. Preliminaries	8
3. Ioffe's theorem	12
4. Slit mappings are almost rigid	21
5. The modulus of extension	31
6. The blob and its boundary	37
7. Extremal length of partial measured foliations	41
8. The blob is semi-smooth	45
9. The blob is a disk	52
10. The deformation retraction	58
11. The remaining cases	64
References	67

2. PRELIMINARIES

Ideal boundary and punctures. A Riemann surface is *hyperbolic* if its universal covering space is conformally isomorphic to the unit disk \mathbb{D} . The only non-hyperbolic Riemann surfaces are $\widehat{\mathbb{C}}$, \mathbb{C} , $\mathbb{C} \setminus 0$ and complex tori, where $\widehat{\mathbb{C}}$ is the Riemann sphere. A hyperbolic surface X can be regarded as the quotient of its universal covering space \mathbb{D} by its group of deck transformations Γ . The *limit set* Λ_Γ is the set of accumulation points in $\partial\mathbb{D}$ of the Γ -orbit of any point $z \in \mathbb{D}$, and the *set of discontinuity* is $\Omega_\Gamma = \partial\mathbb{D} \setminus \Lambda_\Gamma$. The *ideal boundary* of X is $\partial X = \Omega_\Gamma / \Gamma$. The union $X \cup \partial X = (\mathbb{D} \cup \Omega_\Gamma) / \Gamma$ is naturally a bordered Riemann surface, since Γ acts properly discontinuously and analytically on $\mathbb{D} \cup \Omega_\Gamma$. If X is a finite hyperbolic surface, then ∂X has finitely many connected components, each homeomorphic to a circle.

A *puncture* in a Riemann surface X is an end for which there is a proper (preimages of compact sets are compact) conformal embedding $\mathbb{D} \setminus 0 \rightarrow X$. For example, \mathbb{C} has one puncture at infinity and $\mathbb{C} \setminus 0$ has two punctures. Every puncture can be *filled*, meaning that one can add the missing point and extend the complex structure there. The set of punctures of X is denoted by \dot{X} and the Riemann surface obtained by filling the punctures is denoted by $X \cup \dot{X}$. For hyperbolic surfaces, punctures are the same as a cusps, or ends with parabolic monodromy.

Given a finite hyperbolic surface X , we write $\widehat{X} = X \cup \partial X \cup \dot{X}$ for the compact bordered Riemann surface obtained after adding the ideal boundary and filling the punctures. Suppose that ∂X is non-empty. Then if we take two copies of \widehat{X} —the second with reversed orientation—and glue them along ∂X with the identity, we get a closed Riemann surface called the *double* of \widehat{X} . Because of this construction, many theorems about closed Riemann surfaces with finitely many points removed are also true for finite Riemann surfaces. We will state and use such theorems without further remarks.

Montel's theorem. The simplest version of Montel's theorem says that the set of all holomorphic maps from \mathbb{D} to \mathbb{D} is compact. This implies a similar result for holomorphic maps between arbitrary hyperbolic surfaces. A sequence of maps $f_n : X \rightarrow Y$ between Riemann surfaces *diverges locally uniformly* if for every compact sets $K \subset X$ and $L \subset Y$, the sets $f_n(K)$ and L are disjoint for all large enough n . A set \mathcal{F} of maps between two Riemann surfaces X and Y is *normal* if every sequence in \mathcal{F} admits either a locally uniformly convergent subsequence or a locally uniformly divergent subsequence.

Theorem 2.1 (Montel). *If X and Y are hyperbolic surfaces, then every set of holomorphic maps from X to Y is normal.*

See [28, p.34]. Note that the limit of a convergent sequence of holomorphic maps is holomorphic. If every map in the sequence is injective, then the limit is either injective or constant. If every map in the sequence is locally injective, then the limit is either locally injective or constant.

Quasiconformal maps. A K -*quasiconformal map* between Riemann surfaces is a homeomorphism f such that in charts, its first partial derivatives in the distributional sense are locally in L^2 and the formal matrix df of partial derivatives satisfies the inequality $\|df\|^2 \leq K \det(df)$ almost everywhere. For each point z , the real linear map $d_z f$ sends circles in the tangent plane at z to ellipses of eccentricity $\|d_z f\|^2 / \det(d_z f)$ in the tangent plane at $f(z)$, and this ratio is called the *pointwise dilatation* of f at z . The *dilatation* of f , denoted $\text{Dil}(f)$, is the smallest $K \geq 1$ for which f is K -quasiconformal. This is the same as the essential supremum of the pointwise dilatation of f .

A *Beltrami form* on a Riemann surface X is a map $\mu : TX \rightarrow \mathbb{C}$ such that $\mu(\lambda \mathbf{v}) = (\bar{\lambda}/\lambda)\mu(\mathbf{v})$ for every $\mathbf{v} \in TX$ and every $\lambda \in \mathbb{C} \setminus 0$. In charts, the *Wirtinger derivatives* of a locally quasiconformal map f are

$$\partial f = \frac{1}{2} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right) \quad \text{and} \quad \bar{\partial} f = \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right).$$

The ratio $\bar{\partial} f / \partial f$ is naturally a Beltrami form, and is sometimes called the *Beltrami coefficient* of f . The Beltrami coefficient of f encodes the field of ellipses in TX which df sends to circles.

The measurable Riemann mapping theorem says that every measurable ellipse field with bounded eccentricity is the Beltrami coefficient of a quasiconformal map.

Theorem 2.2 (Morrey, Ahlfors–Bers). *Let X be a Riemann surface and let μ be a measurable Beltrami form on X such that $\|\mu\|_\infty < 1$. Then there exists a Riemann surface Y and a quasiconformal map $f : X \rightarrow Y$ such that $\bar{\partial} f / \partial f = \mu$ almost everywhere. The surface and the map are unique up to conformal isomorphism.*

An important consequence is the following factorization principle. Suppose that $f : X \rightarrow Y$ is locally quasiconformal (or quasiregular). Then $f = F \circ \varphi$ where $\varphi : X \rightarrow X'$ is a quasiconformal homeomorphism and $F : X' \rightarrow Y$ is holomorphic. Indeed, we can take φ to be the solution of the Beltrami equation with $\mu = \bar{\partial} f / \partial f$ and let $F = f \circ \varphi^{-1}$.

Montel’s theorem generalizes to suitably normalized K -quasiconformal maps from the disk to itself.

Theorem 2.3. *For every $K \geq 1$, the space of all K -quasiconformal maps from \mathbb{D} to \mathbb{D} fixing 0 is compact.*

Lastly, we will use the fact that isolated points are removable singularities for quasiconformal maps. The reader may consult [2] for background on quasiconformal maps.

Quadratic differentials. A *quadratic differential* on a Riemann surface X is a map $\varphi : TX \rightarrow \widehat{\mathbb{C}}$ such that $\varphi(\lambda \mathbf{v}) = \lambda^2 \varphi(\mathbf{v})$ for every $\mathbf{v} \in TX$ and every $\lambda \in \mathbb{C} \setminus 0$. A quadratic differential on X is *holomorphic* (resp. *meromorphic*) if for every open set $U \subset X$, and every holomorphic vector field $\mathbf{v} : U \rightarrow TU$, the function $\varphi \circ \mathbf{v} : U \rightarrow \mathbb{C}$ is holomorphic (resp. meromorphic). All

quadratic differentials in this paper will be holomorphic or meromorphic. The *pull-back* $f^*\varphi$ of a quadratic differential φ by a holomorphic map f is defined in the usual way by the formula $f^*\varphi(\mathbf{v}) = \varphi(df(\mathbf{v}))$.

A vector $\mathbf{v} \in TX$ is *horizontal* (resp. *vertical*) for φ if $\varphi(\mathbf{v}) > 0$ (resp. $\varphi(\mathbf{v}) < 0$). A smooth arc $\gamma : I \rightarrow X$ is *horizontal* (resp. *vertical*) if $\gamma'(t)$ is horizontal (resp. vertical) for every $t \in I$. The absolute value $|\varphi|$ is an area form, and its integral $\|\varphi\| = \int_X |\varphi|$ is the *norm* of φ . For a finite Riemann surface X , we denote by $\mathcal{Q}(X)$ the set of all integrable holomorphic quadratic differentials φ on X which extend analytically to the ideal boundary of X , and such that $\varphi(\mathbf{v}) \in \mathbb{R}$ for every vector \mathbf{v} tangent to ∂X . Every $\varphi \in \mathcal{Q}(X)$ extends to a meromorphic quadratic differential on \hat{X} with at most simple poles on \hat{X} . The set $\mathcal{Q}^+(X)$ is similarly defined, but with the additional requirements that $\varphi \geq 0$ along ∂X and that φ is not identically zero. The set $\mathcal{Q}(X)$ is a real vector space inside of which $\mathcal{Q}^+(X)$ forms a convex cone.

For every simply connected open set $U \subset X$ where a quadratic differential φ does not have any zero or pole, there exists a locally injective holomorphic map $z : U \rightarrow \mathbb{C}$ such that $\varphi = dz^2$. The map z is unique up to translation and sign and is called a *natural coordinate* when it is injective. If $\varphi \in \mathcal{Q}(X)$, then the atlas of natural coordinates for φ is a half-translation structure on X in the sense of the introduction. Such a structure induces a flat geometry with cone points on X which we will discuss in Sections 3 and 4. The standard reference for this material is [37].

Teichmüller's theorem. A *Teichmüller map* between finite Riemann surfaces X and Y is a homeomorphism $f : X \rightarrow Y$ such that there exists a constant $K > 1$ and non-zero $\varphi \in \mathcal{Q}(X)$ and $\psi \in \mathcal{Q}(Y)$ such that f is locally of the form $x + iy \mapsto Kx + iy$ (up to sign and translation) in natural coordinates. Such a map is K -quasiconformal with constant pointwise dilatation.

The following theorems of Teichmüller are central in the theory of deformations of Riemann surfaces.

Theorem 2.4 (Teichmüller's existence theorem). *Let h be a quasiconformal map between finite Riemann surfaces. If there is no conformal isomorphism homotopic to h , then there is a Teichmüller map homotopic to h .*

Theorem 2.5 (Teichmüller's uniqueness theorem). *Let f be a Teichmüller map of dilatation K between finite Riemann surfaces that are not annuli nor tori. If g is a K -quasiconformal map homotopic to f , then $g = f$.*

For annuli and tori, Teichmüller maps are unique up to conformal automorphisms homotopic to the identity. Teichmüller's theorem is usually stated and proved for closed Riemann surfaces, but the general case follows from the closed case by doubling across the ideal boundary and by taking a branched cover of degree 2 or 4 ramified at the punctures [1].

Reduced Teichmüller spaces. Let S be a finite Riemann surface. The *reduced Teichmüller space* $\mathcal{T}^\#(S)$ is defined as the set of pairs (X, f) where

X is a finite Riemann surface and $f : S \rightarrow X$ is a quasiconformal map, modulo the equivalence relation $(X, f) \sim (Y, g)$ if and only if $g \circ f^{-1}$ is homotopic to a conformal isomorphism. The equivalence class of (X, f) is denoted $[X, f]$, or just X when the *marking* f is implicit. The *Teichmüller distance* between two points of $\mathcal{T}^\#(S)$ is defined as

$$d([X, f], [Y, g]) = \frac{1}{2} \inf \log \text{Dil}(h)$$

where the infimum is taken over all quasiconformal maps h homotopic to $g \circ f^{-1}$. By Teichmüller's theorem, the infimum is realized by a (usually unique) quasiconformal map h which is either conformal or a Teichmüller map.

The space $\mathcal{T}^\#(S)$ is a contractible real-analytic manifold of finite dimension. Let $\mathcal{M}(X)$ denote the space of essentially bounded Beltrami forms on $X \in \mathcal{T}^\#(S)$. By the measurable Riemann mapping theorem, the tangent space to $\mathcal{T}^\#(S)$ at X can be identified with the quotient of $\mathcal{M}(X)$ by its subspace $\mathcal{M}_0(X)$ of infinitesimally trivial deformations. There is a natural pairing between $\mathcal{M}(X)$ and $\mathcal{Q}(X)$ given by

$$\langle \mu, \varphi \rangle = \text{Re} \int_X \mu \varphi,$$

and it turns out that $\mathcal{M}_0(X) = \mathcal{Q}(X)^\perp$ with respect to this pairing. Therefore the tangent and cotangent spaces to $\mathcal{T}^\#(S)$ at X are isomorphic to $\mathcal{M}(X)/\mathcal{Q}(X)^\perp$ and $\mathcal{Q}(X)$ respectively. See [7] and [8] for more details.

Homotopies. If two maps are homotopic, then they induce the same homomorphism between fundamental groups, up to conjugation. The converse also holds under appropriate conditions [2, p.60] [3, §6].

Lemma 2.6 (Ahlfors). *Let X be a space which has a universal cover, let Y be a metric space whose universal cover is a uniquely geodesic space in which geodesics depend continuously on endpoints, and let $f_0, f_1 : X \rightarrow Y$ be continuous maps. Suppose that for some $x \in X$ the induced homomorphisms $\pi_1(f_j) : \pi_1(X, x) \rightarrow \pi_1(Y, f_j(x))$ agree up to conjugation by a path between $f_0(x)$ and $f_1(x)$. Then f_0 and f_1 are homotopic.*

Proof. Let \tilde{X} and \tilde{Y} be the universal covers of X and Y , and let α be a path connecting $f_0(x)$ to $f_1(x)$ which conjugates the homomorphisms $\pi_1(f_0)$ and $\pi_1(f_1)$. Given a lift $\tilde{f}_0 : \tilde{X} \rightarrow \tilde{Y}$, the path α allows us to lift f_1 in such a way that \tilde{f}_0 and \tilde{f}_1 are equivariant with respect to the same homomorphism of deck groups. The homotopy from \tilde{f}_0 to \tilde{f}_1 sending $(x, t) \in \tilde{X} \times [0, 1]$ to the point at proportion t along the geodesic from $\tilde{f}_0(x)$ to $\tilde{f}_1(x)$ in \tilde{Y} is continuous and equivariant, so it descends to a homotopy from f_0 to f_1 . \square

This is also true if X is a CW-complex and Y is a $K(\pi, 1)$ [18, p.90]. If X and Y are hyperbolic surfaces, then either hypotheses are satisfied. The

most useful consequence for us is that homotopy classes of maps from a finite Riemann surface to a hyperbolic surface are closed.

Corollary 2.7. *Let X and Y be hyperbolic surfaces, with X finite, and let $f_n, f : X \rightarrow Y$ be continuous maps such that $f_n \rightarrow f$. Then f_n is homotopic to f for all large enough n .*

Proof. Let β_1, \dots, β_k be loops based at $x \in X$ which generate $\pi_1(X, x)$ and let V be a simply connected neighborhood $f(x)$ in Y .

Let n be large enough so that $f_n(x) \in V$ and $f_n(\beta_j)$ is freely homotopic to $f(\beta_j)$ for every $j \in \{1, \dots, k\}$. If α is any path between $f_n(x)$ and $f(x)$ in V , then $\alpha * f(\beta_j) * \bar{\alpha}$ is homotopic to $f_n(\beta_j)$ for every j . By the previous lemma, f_n is homotopic to f . \square

In Teichmüller theory, one often goes back and forth between punctures and marked points as convenient. This passage is justified by the removability of isolated singularities for quasiconformal maps and the following elementary lemma whose proof is left to the reader.

Lemma 2.8. *Let X and Y be Riemann surfaces, let $f_0, f_1 : X \rightarrow Y$ be continuous maps, and let $x \in X$ and $y \in Y$. Suppose that $f_j^{-1}(y) = \{x\}$ for $j = 0, 1$ and suppose the restrictions $f_0^*, f_1^* : X \setminus x \rightarrow Y \setminus y$ are homotopic. Then there exists a homotopy $H : X \times [0, 1] \rightarrow Y$ from f_0 to f_1 such that $H^{-1}(y) = \{x\} \times [0, 1]$.*

The analogous statement for finite sets of points holds as well.

3. IOFFE'S THEOREM

The goal of this section is to characterize quasiconformal embeddings that have minimal dilatation in their homotopy class. We start with a compactness lemma that guarantees the existence of such *extremal* quasiconformal embeddings. From this section to Section 10, **X and Y are assumed to be finite hyperbolic surfaces.**

Definition 3.1. An embedding between Riemann surfaces is *parabolic* if the image of the induced homomorphism on fundamental groups is cyclic and generated by a loop around a puncture.

Recall also that h is *trivial* or *cyclic* if the image of $\pi_1(h)$ is trivial or infinite cyclic respectively.

Lemma 3.2 (Ioffe). *Let $h : X \rightarrow Y$ be a quasiconformal embedding and let $K \geq 1$. The space of all K -quasiconformal embeddings homotopic to h is compact if and only if h is neither trivial nor parabolic.*

Proof. Let $\pi_X : \mathbb{D} \rightarrow X$ and $\pi_Y : \mathbb{D} \rightarrow Y$ be universal covering maps with respective deck groups $\text{Deck}(X)$ and $\text{Deck}(Y)$. The map h induces a homomorphism $\Theta : \text{Deck}(X) \rightarrow \text{Deck}(Y)$ once we fix basepoints. If $f : X \rightarrow Y$ is a K -quasiconformal embedding homotopic to h , then it lifts to a locally

K -quasiconformal map $\tilde{f} : \mathbb{D} \rightarrow \mathbb{D}$ satisfying $\tilde{f}(\alpha \cdot z) = \Theta(\alpha) \cdot \tilde{f}(z)$ for every $z \in \mathbb{D}$ and every $\alpha \in \text{Deck}(X)$. The lift \tilde{f} is only defined up to composition by elements of the two deck groups. The condition that f be injective is equivalent to the condition $\pi_Y(\tilde{f}(z)) = \pi_Y(\tilde{f}(w)) \Rightarrow \pi_X(z) = \pi_X(w)$.

There is a unique factorization $\tilde{f} = F \circ \varphi$ where $\varphi : \mathbb{D} \rightarrow \mathbb{D}$ is a K -quasiconformal homeomorphism fixing 0 and 1, and where $F : \mathbb{D} \rightarrow \mathbb{D}$ is holomorphic and locally injective. As mentioned earlier, the space of K -quasiconformal homeomorphisms of \mathbb{D} fixing the origin and the space of holomorphic maps from \mathbb{D} to $\overline{\mathbb{D}}$ are both compact. Thus, given a sequence of K -quasiconformal embeddings $f_n : X \rightarrow Y$ and corresponding lifts $\tilde{f}_n = F_n \circ \varphi_n$, we can pass to a subsequence such that $F_n \rightarrow F$ and $\varphi_n \rightarrow \varphi$ and hence $\tilde{f}_n \rightarrow \tilde{f} = F \circ \varphi$.

Suppose that each f_n is homotopic to h and that the latter is neither trivial nor cyclic. Then the limit \tilde{f} is not constant. Indeed, if there is a constant $w_0 \in \overline{\mathbb{D}}$ such that $\tilde{f}(z) = w_0$ for every $z \in \mathbb{D}$, then

$$w_0 = \tilde{f}(\alpha \cdot z) = \lim_{n \rightarrow \infty} \tilde{f}_n(\alpha \cdot z) = \lim_{n \rightarrow \infty} \Theta(\alpha) \cdot \tilde{f}_n(z) = \Theta(\alpha) \cdot \tilde{f}(z) = \Theta(\alpha) \cdot w_0$$

for every $\alpha \in \text{Deck}(X)$. However, a Fuchsian group which fixes a point in $\overline{\mathbb{D}}$ is cyclic, contradicting the assumption on h . In particular, the holomorphic function φ is not constant, thus has image in \mathbb{D} , and by Hurwitz's theorem it is locally injective. Therefore $\tilde{f} = F \circ \varphi$ is locally K -quasiconformal. Moreover, the equality $\tilde{f}(\alpha \cdot z) = \Theta(\alpha) \cdot \tilde{f}(z)$ for every $z \in \mathbb{D}$ and every $\alpha \in \text{Deck}(X)$ implies that \tilde{f} descends to a locally K -quasiconformal map $f : X \rightarrow Y$.

We have to show that f is injective. If $f(\pi_X(z)) = f(\pi_X(w))$, then $\pi_Y(\tilde{f}(z)) = \pi_Y(\tilde{f}(w))$. Since $\tilde{f}_n \rightarrow \tilde{f}$ and since these maps are open, we can find a sequence z_n converging to z and a sequence w_n converging to w such that $\tilde{f}_n(z_n) = \tilde{f}(z)$ and $\tilde{f}_n(w_n) = \tilde{f}(w)$ for all n large enough. Then $\pi_Y(\tilde{f}_n(z_n)) = \pi_Y(\tilde{f}_n(w_n))$, so that $\pi_X(z_n) = \pi_X(w_n)$, and thus $\pi_X(z) = \pi_X(w)$. Lastly, f is homotopic to h because it is a limit of maps which are.

Suppose now that the image of Θ is a cyclic group generated by a hyperbolic element $\beta \in \text{Deck}(Y)$. Let $z_0 \in \mathbb{D}$, and let $D \subset \mathbb{D}$ be a fundamental domain for β whose closure is disjoint from the fixed points of β . By applying an appropriate power of β , we can assume that the lift \tilde{f}_n is such that $\tilde{f}_n(z_0) \in D$. After such a normalization, the limit $\tilde{f}(z_0)$ belongs to the closure \overline{D} , so it is not one of the fixed points of β . It follows that \tilde{f} is not constant. Then the above argument applies verbatim.

Conversely, if h is trivial then its image is contained in a disk in Y . In other words, h can be written as $h = F \circ \varphi$ where $F : \mathbb{D} \rightarrow Y$ is a conformal embedding and $\varphi : X \rightarrow \mathbb{D}$ is a K -quasiconformal embedding. Consider the sequence $f_n = F_n \circ \varphi$ where $F_n(z) = F(z/n)$. Each f_n is a K -quasiconformal embedding isotopic to h , but the sequence converges to a constant map. Similarly, if h is parabolic, then we can form a sequence of

isotopic K -quasiconformal embeddings which diverges to the corresponding puncture.

□

The following corollary is immediate.

Corollary 3.3. *Let $h : X \rightarrow Y$ be a non-trivial and non-parabolic quasiconformal embedding. There exists a quasiconformal embedding $f : X \rightarrow Y$ homotopic to h with minimal dilatation.*

We present Ioffe's theorem in two parts. The first part says that every extremal quasiconformal embedding is a Teichmüller embedding.

Definition 3.4. A *Teichmüller embedding* of dilatation $K \geq 1$ is an injective continuous map $f : X \rightarrow Y$ for which there exist quadratic differentials $\varphi \in \mathcal{Q}^+(X)$ and $\psi \in \mathcal{Q}^+(Y)$ such that f has the form $x + iy \mapsto Kx + iy$ in natural coordinates and such that $Y \setminus f(X)$ is an analytic graph all of whose edges are horizontal for ψ . We say that φ and ψ are *initial* and *terminal* quadratic differentials for f . A *slit mapping* is a conformal Teichmüller embedding, i.e. one with $K = 1$.

If f is a Teichmüller embedding of dilatation K with respect to φ and ψ , then $\bar{\partial}f/\partial f = k\bar{\varphi}/|\varphi|$ and $\bar{\partial}(f^{-1})/\partial(f^{-1}) = -k\bar{\psi}/|\psi|$ where $k = \frac{K-1}{K+1}$.

Theorem 3.5 (Ioffe). *Let $f : X \rightarrow Y$ be a quasiconformal embedding with minimal dilatation in its homotopy class. If f is not conformal, then it is a Teichmüller embedding.*

Proof. Suppose that f is not conformal and let

$$\mu := \begin{cases} \bar{\partial}(f^{-1})/\partial(f^{-1}) & \text{on } f(X) \\ 0 & \text{on } Y \setminus f(X). \end{cases}$$

Let $F : Y \rightarrow Y_\mu$ be the solution to the Beltrami equation $\bar{\partial}F/\partial F = \mu$. By construction we have $\bar{\partial}(F \circ f)/\partial(F \circ f) = 0$ so that $F \circ f$ is a conformal embedding. Let $G : Y_\mu \rightarrow Y$ be a quasiconformal map homotopic to F^{-1} with minimal dilatation. By Teichmüller's theorem, either G is conformal or a Teichmüller map. The composition $G \circ F \circ f : X \rightarrow Y$ is a quasiconformal embedding homotopic to f , so that

$$\text{Dil}(f) \leq \text{Dil}(G \circ F \circ f) \leq \text{Dil}(G) \leq \text{Dil}(F^{-1}) = \text{Dil}(f).$$

Thus all the terms in this chain are equal. The equality $\text{Dil}(G) = \text{Dil}(F^{-1})$ implies that F has minimal dilatation in its homotopy class. Since F is not conformal, it is a Teichmüller map. This means that there is a non-zero $\psi \in \mathcal{Q}(Y)$ and a constant $k \in (0, 1)$ such that $\mu = -k\bar{\psi}/|\psi|$ almost everywhere. In particular, $Y \setminus f(X)$ has measure zero and f has constant pointwise dilatation.

Thought of as a homeomorphism from X to $f(X)$, the map f has minimal dilatation in its homotopy class and is thus a Teichmüller map. Let φ and ω

be its initial and terminal quadratic differentials. Since $F \circ f$ is conformal, the directions of maximal stretching for F and f^{-1} must be perpendicular, which means that $\psi = c\omega$ on $f(X)$ for some positive constant c , which we may assume is equal to 1 by rescaling.

We have to show that $\varphi \in \mathcal{Q}^+(X)$. If not, then $\varphi < 0$ along some segment $I \subset \partial X$. We will explicitly construct a quasiconformal embedding \tilde{f} from X to $f(X)$ with pointwise dilatation smaller than f near I . We may work in a natural coordinate chart for φ in which I is equal to the vertical segment $[-i, i]$ in the plane and X is to the right of I . There is also a natural chart for ω in which f takes the form $x + iy \mapsto Kx + iy$. Let Δ be the isocles triangle with base $[-i, i]$ and apex $\delta > 0$. Consider the map $L : \Delta \rightarrow \Delta$




FIGURE 1. Reducing the pointwise dilatation of an embedding near a vertical boundary arc.

which is affine on the upper and lower halves of Δ , fixes all three vertices of Δ , and sends the midpoint of I to $(1 - 1/K)\delta$. Extend L to be the identity on $X \setminus \Delta$ and let $\tilde{f} = f \circ L$. The linear part of \tilde{f} on the lower half of Δ is equal to

$$\begin{pmatrix} 1 & (K-1)\delta \\ 0 & 1 \end{pmatrix}$$

and the dilatation of this matrix tends to 1 as $\delta \rightarrow 0$. Therefore, if δ is small enough, then the embedding \tilde{f} has strictly smaller pointwise dilatation than f on Δ . Moreover the global dilatation of \tilde{f} is the same as f , so that \tilde{f} also has minimal dilatation in its homotopy class. By the first paragraph of the proof, the pointwise dilatation of \tilde{f} must be constant. This is a contradiction, and hence $\varphi \in \mathcal{Q}^+(X)$.

It remains to show that $f(X)$ is the complement of a graph which is horizontal with respect to ψ . Recall that \hat{X} is the compactification of X obtained by adding its ideal boundary and filling its punctures. The path metrics induced by $|\varphi|^{1/2}$ and $|\psi|^{1/2}$ extend to complete metrics on \hat{X} and \hat{Y} . Since $f : X \rightarrow Y$ is K -Lipschitz with respect to those metrics, it extends to a K -Lipschitz map $\hat{f} : \hat{X} \rightarrow \hat{Y}$. Moreover, \hat{f} is surjective since $Y \setminus f(X)$ has measure zero and hence empty interior. Let I be the closure of a connected component of $\partial X \setminus \{\text{zeros of } \varphi\}$. There is a sequence $\{I_n\}$ of arcs in X which are horizontal for φ and converge uniformly to I . Since the image arcs $f(I_n)$

are all horizontal for ψ , they can only accumulate onto horizontal arcs, and thus $\widehat{f}(I)$ is horizontal. Therefore, the complement $\widehat{Y} \setminus f(X) = \widehat{f}(\partial X \cup \dot{X})$ is a union of finitely many points and horizontal arcs for ψ . In particular, the ideal boundary ∂Y is horizontal for ψ so that $\psi \in \mathcal{Q}^+(Y)$. \square

In the last paragraph of the proof we actually showed the following useful criterion.

Lemma 3.6. *Let $\varphi \in \mathcal{Q}^+(X)$, let $\psi \in \mathcal{Q}(Y)$, and let $f : X \rightarrow Y$ be an embedding which is locally of the form $x + iy \mapsto Kx + iy$ in natural coordinates. If $f(X)$ is dense in Y , then $\psi \in \mathcal{Q}^+(Y)$ and f is a Teichmüller embedding.*

The second part of Ioffe's theorem says that every Teichmüller embedding is extremal. The proof is very similar to the proofs of Teichmüller's uniqueness theorem given in [3] and [11, Chapter 11].

Theorem 3.7 (Ioffe). *Let $f : X \rightarrow Y$ be a Teichmüller embedding of dilatation K with terminal quadratic differential ψ , and let $g : X \rightarrow Y$ be a K -quasiconformal embedding homotopic to f . Then g is a Teichmüller embedding of dilatation K with terminal quadratic differential ψ , and the map $g \circ f^{-1} : f(X) \rightarrow g(X)$ is locally a translation in natural coordinates.*

Proof. Let φ be the initial quadratic differential of f . By rescaling, we may assume that $\|\varphi\| = 1$. This implies that $\|\psi\| = K$ since f multiplies area by a factor K and $f(X)$ has full measure in Y . We equip \widehat{X} and \widehat{Y} with the complete conformal metrics $|\varphi|^{1/2}$ and $|\psi|^{1/2}$. Since f is K -Lipschitz with respect to those metrics, it extends to a K -Lipschitz map $\widehat{f} : \widehat{X} \rightarrow \widehat{Y}$.

We may assume that \widehat{f} maps \dot{X} into \dot{Y} . Indeed, we can fill all the punctures of Y that are not poles of ψ , and then fill all the punctures of X that do not map to poles of ψ . The map f extended to those points is still a Teichmüller embedding.

Let $M > 1$. For every $n \in \mathbb{N}$, let $G_n : \widehat{X} \rightarrow \widehat{X}$ be a smooth M -quasiconformal and M -Lipschitz embedding such that

- $G_n(\dot{X}) = \dot{X}$;
- G_n is homotopic to the identity rel \dot{X} ;
- $G_n(\widehat{X})$ is contained in $X \cup \dot{X}$;
- $G_n \rightarrow \text{id}$ uniformly in the C^1 norm as $n \rightarrow \infty$.

Such maps are not difficult to construct. Given a smooth vector field on \widehat{X} pointing inwards on ∂X and vanishing on \dot{X} , we can let G_n be the corresponding flow at small enough time t_n . Then let $g_n = g \circ G_n : \widehat{X} \rightarrow \widehat{Y}$, which is a KM -quasiconformal embedding.

Let $H : \widehat{X} \times [0, 1] \rightarrow \widehat{Y}$ be a homotopy from \widehat{f} to g_n which is constant on \dot{X} . For every $z \in X$, let $\ell(z) = \inf_{\alpha} \int_{\alpha} |\psi|^{1/2}$ where the infimum is taken over all piecewise smooth paths $\alpha : [0, 1] \rightarrow Y$ that are homotopic to the path $t \mapsto H(z, t)$ rel endpoints. If \dot{Y} is empty, then the infimum $\ell(z)$ is realized

by a unique geodesic. However, in general the infimum need not be realized since the restriction of $|\psi|^{1/2}$ to Y is not complete. In any case, it is easy to see that ℓ is continuous on X and extends continuously to \hat{X} . Since \hat{X} is compact, there exists a constant B such that $\ell(z) < B$ for every $z \in \hat{X}$.

Let γ be a horizontal arc of length L in X . Since f is a Teichmüller embedding of dilatation K , it sends γ to a horizontal arc of length KL in Y . Let z_0 and z_1 be the endpoints of γ . We can obtain a path homotopic to $f(\gamma)$ by taking the concatenation of a piecewise smooth path α_0 of length at most B homotopic to $t \mapsto H(z_0, t)$, the image $g_n(\gamma)$, and a piecewise smooth path α_1 of length at most B homotopic to $t \mapsto H(z_1, 1 - t)$. Since horizontal arcs minimize horizontal travel among all homotopic paths [37, p.76], we have

$$\begin{aligned} KL &= \int_{f(\gamma)} |\operatorname{Re} \sqrt{\psi}| \leq \int_{\alpha_0} |\operatorname{Re} \sqrt{\psi}| + \int_{g_n(\gamma)} |\operatorname{Re} \sqrt{\psi}| + \int_{\alpha_1} |\operatorname{Re} \sqrt{\psi}| \\ &\leq 2B + \int_{g_n(\gamma)} |\operatorname{Re} \sqrt{\psi}| \end{aligned}$$

Let dg_n denote the matrix of partial derivatives of g_n with respect to natural coordinates³ and $(dg_n)_{1,1}$ its first entry. If z and ζ are natural coordinates for φ and ψ , then $(dg_n)_{1,1} = \operatorname{Re}(\partial(\zeta \circ g_n \circ z^{-1})/\partial x)$. If g_n is absolutely continuous on γ , then we have

$$\int_{\gamma} |(dg_n)_{1,1}| \cdot |\varphi|^{1/2} = \int_{g_n(\gamma)} |\operatorname{Re} \sqrt{\psi}| \geq KL - 2B.$$

Remove from X all trajectories that go through a puncture of X or a zero of φ and denote the resulting full measure subset by U . For every $z \in U$, there is a unique (possibly closed) bi-infinite horizontal trajectory through z . For every $L > 0$ and every $z \in U$, let γ_z^L be the horizontal arc of length L centered at z . Since g_n is quasiconformal, it is absolutely continuous on almost every horizontal trajectory. Upon applying Fubini's theorem, we find

$$\int_U |(dg_n)_{1,1}| \cdot |\varphi| = \int_U \left(\frac{1}{L} \int_{\gamma_z^L} |(dg_n)_{1,1}| \cdot |\varphi|^{1/2} \right) \cdot |\varphi| \geq \left(K - \frac{2B}{L} \right) \int_U |\varphi|.$$

Letting $L \rightarrow \infty$, we obtain $\int_X |(dg_n)_{1,1}| \cdot |\varphi| \geq K \int_X |\varphi| = K$, since U has full measure in X .

We claim that $\int_X |(dg)_{1,1}| \cdot |\varphi| = \lim_{n \rightarrow \infty} \int_X |(dg_n)_{1,1}| \cdot |\varphi|$ and hence that $\int_X |(dg)_{1,1}| \cdot |\varphi| \geq K$. This is a consequence of the Vitali convergence theorem [33, p.94]. In order to apply the theorem, we need to check that the functions $|(dg_n)_{1,1}|$ are uniformly integrable. First observe that $\int_X \det(dg) \cdot |\varphi| = \int_{g(X)} |\psi| \leq K$, so that $\det(dg)$ is integrable. It follows that for every $\varepsilon > 0$, there exists a $\delta > 0$ such that if $A \subset X$ is measurable and $\int_A |\varphi| < \delta$,

³The matrix is only defined up to sign, but no matter.

then $\int_A \det(dg) \cdot |\varphi| < \varepsilon$. Now if $\int_A |\varphi| < \delta/M^2$, then $\int_{G_n(A)} |\varphi| < \delta$ since G_n is M -Lipschitz. By the Cauchy-Schwarz inequality we have

$$\begin{aligned} \left(\int_A |(\mathrm{d}g_n)_{1,1}| \cdot |\varphi| \right)^2 &\leq \int_A |(\mathrm{d}g_n)_{1,1}|^2 \cdot |\varphi| \leq \int_A \|\mathrm{d}g_n\|^2 \cdot |\varphi| \\ &\leq KM \int_A \det(\mathrm{d}g_n) \cdot |\varphi| \\ &= KM \int_A \det(\mathrm{d}_{G_n(z)}g) \det(\mathrm{d}_z G_n) \cdot |\varphi| \\ &\leq KM \int_{G_n(A)} \det(\mathrm{d}g) \cdot |\varphi| < KM\varepsilon, \end{aligned}$$

which shows uniform integrability and proves the claim.

Applying Cauchy-Schwarz to the inequality $K \leq \int_X |(\mathrm{d}g)_{1,1}| \cdot |\varphi|$ yields

$$\begin{aligned} K^2 &\stackrel{(a)}{\leq} \left(\int_X |(\mathrm{d}g)_{1,1}| \cdot |\varphi| \right)^2 \stackrel{(b)}{\leq} \int_X |(\mathrm{d}g)_{1,1}|^2 \cdot |\varphi| \\ &\stackrel{(c)}{\leq} \int_X \|\mathrm{d}g\|^2 \cdot |\varphi| \leq K \int_X \det(\mathrm{d}g) \cdot |\varphi| \\ &= K \int_{g(X)} |\psi| \stackrel{(d)}{\leq} K \int_Y |\psi| = K^2. \end{aligned}$$

Since the two ends of this chain of inequalities agree, each intermediate inequality is in fact an equality. Equality in (b) implies that $|(\mathrm{d}g)_{1,1}|$ is equal to a constant almost everywhere on X , and that constant is equal to K by (a). The inequality (c) is based on

$$|(\mathrm{d}g)_{1,1}|^2 \leq |(\mathrm{d}g)_{1,1}|^2 + |(\mathrm{d}g)_{1,2}|^2 = |\mathrm{d}g\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right)|^2 \leq \sup_{\|\mathbf{v}\|=1} |\mathrm{d}g(\mathbf{v})|^2 = \|\mathrm{d}g\|^2.$$

Equality implies that $\mathrm{d}g\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right) = \pm\left(\begin{smallmatrix} K \\ 0 \end{smallmatrix}\right)$. Moreover, since $\mathrm{d}g$ stretches maximally in the horizontal direction which is preserved, $\mathrm{d}g$ must be diagonal, i.e. $\mathrm{d}g = \pm\left(\begin{smallmatrix} K & 0 \\ 0 & * \end{smallmatrix}\right)$ with $0 < * \leq K$. Then the equality $K^2 = \|\mathrm{d}g\|^2 = K \det(\mathrm{d}g)$ determines that $\mathrm{d}g = \pm\left(\begin{smallmatrix} K & 0 \\ 0 & 1 \end{smallmatrix}\right)$ almost everywhere on X .

Since $\mathrm{d}f = \left(\begin{smallmatrix} K & 0 \\ 0 & 1 \end{smallmatrix}\right)$ up to sign as well, we have $\mathrm{d}(g \circ f^{-1}) = \pm\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ almost everywhere on $f(X)$. The Beltrami coefficient of $g \circ f^{-1}$ is thus equal to 0 almost everywhere on $f(X)$, so that $g \circ f^{-1}$ is conformal and in particular smooth. Since f is smooth except at the zeros of φ , the same holds for g . Therefore the equality $\mathrm{d}g = \pm\left(\begin{smallmatrix} K & 0 \\ 0 & 1 \end{smallmatrix}\right)$ holds everywhere except at the zeros of φ , and g takes the form $x + iy \mapsto \pm(Kx + iy) + c$ in natural coordinates. Equality in (d) means that $g(X)$ has full measure in Y . By Lemma 3.6, g is a Teichmüller embedding with terminal quadratic differential ψ . Finally, the equality $\mathrm{d}(g \circ f^{-1}) = \pm\left(\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ holds everywhere, so that $g \circ f^{-1}$ is a local translation. \square

In the case where X and Y are closed hyperbolic surfaces, we get that $g \circ f^{-1} : Y \rightarrow Y$ is a conformal isomorphism (hence an isometry) homotopic

to the identity. It is well known [3, §12.A] that the identity is the only such map, so that $g = f$. This fact is also used in the proof of Hurwitz's theorem on automorphisms of closed hyperbolic surfaces [11, p.215]. We prove a more general statement for later use.

Lemma 3.8. *Let X be a hyperbolic Riemann surface not isomorphic to \mathbb{D} nor $\mathbb{D} \setminus 0$, and let $h : X \rightarrow X$ be a holomorphic map homotopic to the identity. Then h is equal to the identity unless X is an annulus and h is a rotation.*

Proof. By the Schwarz lemma, h is 1-Lipschitz with respect to the hyperbolic metric. Therefore, if α is a closed geodesic in X , then $h(\alpha)$ is at most as long as α . But geodesics minimize length in their homotopy class, so that $h(\alpha) = \alpha$. In particular, h is an isometry along α .

If X is an annulus, then it contains a unique simple closed geodesic α . We can post-compose h by a rotation r of X so that $r \circ h$ is equal to the identity on α and hence on all of X by the identity principle.

If X is not an annulus, then it contains a closed geodesic α which self-intersects exactly once. Then h fixes this self-intersection point, thus all of α pointwise, and hence all of X pointwise by the identity principle. \square

As pointed out before the proof, this lemma implies the uniqueness of Teichmüller maps between closed hyperbolic surfaces. However, the above argument does not apply to non-surjective Teichmüller embeddings. Indeed, Teichmüller embeddings are not necessarily unique in their homotopy class.

If f and g are homotopic Teichmüller embeddings from X to Y , then the inclusion $f(X) \hookrightarrow Y$ and the map $g \circ f^{-1} : f(X) \rightarrow Y$ are homotopic slit mappings by Theorem 3.7. Therefore, it suffices to study the question of uniqueness for slit mappings. There are two obvious ways for uniqueness to fail:

- if Y is a torus⁴, then we can post-compose f with any automorphism of Y isotopic to the identity;
- if $f(X)$ is contained in an annulus $A \subset Y$, then we can post-compose f with rotations of A .

In [20], Ioffe claims these are the only exceptions, but this is wrong⁵. The next simplest example is as follows. Let Y be a round annulus in the plane with a concentric circular arc removed, and let X be the same annulus but with a slightly longer arc removed. Then we can obviously rotate X inside of Y by some amount. This gives a 1-parameter family of slit mappings between triply connected domains. If the annulus is centered at the origin, then the terminal quadratic differential is $-dz^2/z^2$.

Figure 2 shows this example, but from a different perspective. The slit in Y was opened up and mapped on the outside. This is to illustrate a

⁴This case is not relevant here since we assume Y to be hyperbolic, but nevertheless.

⁵The source of the mistake is [20, Lemma 3.2]. Similarly, [21], [6], and [17] contain minor errors as they build up on the false claim.




FIGURE 2. A 1-parameter family of slit mappings between two pairs of pants. The domain X is the complement of the red slits. The blue arrows indicate the flow by rotations.

point which we have not explained yet. If Y is a finite Riemann surface, $\psi \in \mathcal{Q}^+(Y)$, and X is the complement of a collection of horizontal slits in Y , then $\iota^*\psi$ belongs to $\mathcal{Q}^+(X)$, where $\iota : X \rightarrow Y$ is the inclusion map. The only thing to check is that $\iota^*\psi$ extends analytically to the ideal boundary of X . Near the endpoint of a slit, the process of unfolding the slit to an ideal boundary component is the same as taking a square root. If we pull-back the quadratic differential dz^2 in \mathbb{C} by the square root from $\mathbb{C} \setminus [0, \infty)$ to the upper half-plane, we get the quadratic differential $4z^2 dz^2$. In other words, unfolding a regular point of ψ yields half of a double zero of $\iota^*\psi$ on the ideal boundary. More generally, if the slit ends at a zero of order k , unfolding yields half of a zero of order $2k + 2$. It is perhaps more natural to count the number of prongs: an n -prong singularity transforms into half of a $2n$ -prong singularity (see Figure 3). Clearly ∂X is horizontal for $\iota^*\psi$.




FIGURE 3. Unfolding a slit ending at a simple zero of a quadratic differential yields half of a zero of order 4.

Here is a general method for constructing examples of non-unique slit mappings with arbitrary codomain Y . Let ω be a holomorphic 1-form on Y such that $\omega^2 \in \mathcal{Q}^+(Y)$. If $X \subset Y$ is the complement of enough horizontal slits for ω^2 , then X is not rigid. By “enough” we mean that for every point $y \in \hat{Y}$ which is either a zero of ω or a puncture of Y , and every trajectory γ of ω ending at y in forward time, we must remove a neighborhood of y in γ . Once this is done, then for all small enough $t > 0$ the time- t flow for the vector field $1/\omega$ is conformal and injective on X , and is homotopic to the inclusion map yet different from it.

In the next section, we show that this kind of phenomenon is the worse that can happen: we can always get from any slit mapping to any homotopic

one by flowing horizontally, although the quadratic differential does not have to be globally the square of an abelian differential.

4. SLIT MAPPINGS ARE ALMOST RIGID

We prove that slit mappings are almost rigid in their isotopy class, in the sense that we can get from any one to any other by flowing horizontally.

Theorem 4.1. *Let $f_0, f_1 : X \rightarrow Y$ be homotopic slit mappings with terminal quadratic differential ψ . Then there is an isotopy f_t from f_0 to f_1 through slit mappings with terminal quadratic differential ψ . If $f_0 \neq f_1$, then for every $x \in X$ the path $t \mapsto f_t(x)$ is horizontal for ψ and its length is independent of x .*

The statement is still true if we replace each occurrence of “slit mappings” with “Teichmüller embeddings of dilatation K ”. The general case follows from the case $K = 1$ by Theorem 3.7. The basic idea of the proof of theorem 4.1 is to use the flat metric $|\psi|^{1/2}$ on Y to modify the given homotopy H from f_0 to f_1 into a more geometric homotopy F . Based on the examples of non-rigid slit mappings we saw so far, we want $t \mapsto F(x, t)$ to be the geodesic homotopic to $t \mapsto H(x, t)$ on \hat{Y} . If ψ has simple poles on \dot{Y} then geodesics are not unique, which is problematic. We use the following idea from [1, Section II.3] to get rid of the poles.

Lemma 4.2 (Ahlfors). *Suppose that \dot{Y} is non-empty. There exists a finite hyperbolic surface W and a covering map $c : W \rightarrow Y$ of degree 2 or 4 which extends to a holomorphic map $c : \hat{W} \rightarrow \hat{Y}$ which is branched exactly at the points of $c^{-1}(\dot{Y})$ where it has local degree 2.*

Proof. Let $k = \#\dot{Y}$. Suppose first that k is even. Then we can find disjoint simple smooth arcs in Y joining the points of \dot{Y} in pairs. Take two copies of \hat{Y} , cut each copy along the arcs, glue the left side of each arc in the first copy to the right side of the corresponding arc in the second copy and vice versa, let W' be the resulting bordered Riemann surface, and let W be the interior of W' minus the endpoints of the cuts. Then $\hat{W} = W'$ and forgetting which copy of \hat{Y} a point came from yields the desired degree 2 holomorphic branched cover c . The map c is locally injective except at the endpoints of the cuts where it has local degree 2. The restriction of c to W is thus a covering map.

If k is odd and larger than 1, then we can do the above construction with $(k - 1)$ points. The point left over has two distinct lifts in the resulting double branched cover. We can thus repeat the first construction with these two points, and get a degree 4 cover with degree 2 branched points over \dot{Y} .

Suppose that $k = 1$. Since Y is hyperbolic, either it has positive genus or non-empty ideal boundary. If Y has positive genus, then we can find a non-separating simple smooth loop in Y , cut \hat{Y} along that loop, and glue two copies along the loop. This gives an unbranched degree two cover in which

\dot{Y} has two distinct lifts, so we can perform the first construction with these two lifts. If \widehat{Y} has non-empty boundary, then we can cut \widehat{Y} along a simple smooth arc in Y joining \dot{Y} to the boundary, and glue two copies along the arc. The forgetful map is then a degree two branched cover ramified only over \dot{Y} . \square

Remark. We do not use the orientation double cover of ψ because we do not want the map c to be branched over the zeros of ψ in Y .

Let W and c be as in the previous lemma. If \dot{Y} is empty, then we let $W = Y$ and take c to be the identity map. We then let \mathcal{U} be the universal covering space of \widehat{W} and we let $b : \mathcal{U} \rightarrow \widehat{Y}$ be the composition of $\mathcal{U} \rightarrow \widehat{W}$ with the branched cover $c : \widehat{W} \rightarrow \widehat{Y}$. Another way to say this is: consider \widehat{Y} as an orbifold with the group $\mathbb{Z}/2\mathbb{Z}$ attached to each point of \dot{Y} , and let $b : \mathcal{U} \rightarrow \widehat{Y}$ be the corresponding universal covering orbifold.

To avoid unnecessarily heavy notation, we denote the pull-back $b^*\psi$ by the same letter ψ . The quadratic differential ψ is holomorphic on \mathcal{U} since b is branched of degree 2 over every point of \dot{Y} . The metric $|\psi|^{1/2}$ induces a distance on \mathcal{U} defined by $d(x, y) = \inf_{\alpha} \int_{\alpha} |\psi|^{1/2}$ where the infimum is taken over all piecewise smooth paths α between x and y .

We recall some well-known facts about the geometry of (\mathcal{U}, d) . First, the metric d is complete since the corresponding metric on \widehat{W} is complete. A *minimizing geodesic* in (\mathcal{U}, d) is a shortest path between two points, whereas a *geodesic* is a locally length minimizing path between two points. Since (\mathcal{U}, d) is locally compact and complete, the Hopf–Rinow theorem implies that any two points in \mathcal{U} can be joined by a minimizing geodesic, and that closed bounded sets in \mathcal{U} are compact [4, p.35].

The metric induced by $|\psi|^{1/2}$ is simply the pull-back of the Euclidean metric by natural coordinates. Thus a geodesic in \mathcal{U} travels in straight lines in natural coordinates. Near a k -pronged singularity the metric is isometric to a cone with angle $k\pi$ and a geodesic cannot turn inefficiently when it hits such a cone point. More precisely, a path $\gamma : [0, 1] \rightarrow \mathcal{U}$ is geodesic if and only if

- γ is smooth except perhaps at zeros of ψ ;
- the argument of $\psi(\gamma'(t))$ is locally constant where γ is smooth;
- at a zero of ψ , the cone angle on either side of γ is at least π .

The last point is called the *angle condition*.

A *geodesic polygon* in \mathcal{U} is a concatenation of finitely many geodesics which closes up. Any geodesic polygon in \mathcal{U} is compact and thus the region it bounds contains only finitely many zeros of ψ . Therefore, the Gauss–Bonnet theorem can be applied, and it takes a very simple form in this context.

Lemma 4.3 (Teichmüller). *Let P be a simple geodesic polygon in \mathcal{U} . Then*

$$\sum_{z \in \text{int}(P)} (2\pi - \angle z) + \sum_{z \in \partial P} (\pi - \angle z) = 2\pi.$$

Proof. The Gauss–Bonnet formula says that

$$\int_{\text{int}(P)} K \, dA + \int_{\partial P} k \, ds = 2\pi\chi(P) = 2\pi.$$

The Gaussian curvature K of the metric $|\psi|^{1/2}$ is nil everywhere in the interior of P except at a zero z of ψ , where the curvature is equal to 2π minus the cone angle at z (which is the same as $-\pi$ times the order of the zero). The geodesic curvature k is also zero everywhere along ∂P except at vertices of the polygon or at zeros of ψ . The geodesic curvature at such points is the “turning angle”, or π minus the internal angle. \square

This implies that there is a unique geodesic between any two points of \mathcal{U} . In particular, every geodesic is minimizing.

Lemma 4.4. *Any two points of \mathcal{U} are connected by a unique geodesic.*

Proof. Suppose that $x, y \in \mathcal{U}$ can be joined by two distinct geodesics α and β . Then α and β bound one (or more) simply connected polygon P . The boundary of the polygon is geodesic except for two points p and q where α and β intersect. Thus the turning angle of ∂P is non-positive everywhere except perhaps at p and q by the angle condition. Teichmüller’s lemma gives

$$2\pi = \sum_{z \in \text{int}(P)} (2\pi - \angle z) + \sum_{z \in \partial P \setminus \{p, q\}} (\pi - \angle z) + (2\pi - \angle p - \angle q) < 2\pi,$$

which is a contradiction. \square

The Arzelà–Ascoli theorem and the uniqueness of geodesics together imply that geodesics depend continuously on endpoints [4, p.37].

Lemma 4.5. *The geodesics in \mathcal{U} depend uniformly continuously on endpoints.*

The next theorem is a generalization of Strebel’s “divergence principle”: if two geodesic rays γ_0 and γ_1 are such that the angles they form with the geodesic between $\gamma_0(0)$ and $\gamma_1(0)$ sum to at least π , then the distance between $\gamma_0(t)$ and $\gamma_1(t)$ is non-decreasing. In the original statement [37, p.77], each angle is assumed to be at least $\frac{\pi}{2}$. Our proof makes use of the following immediate corollaries of Teichmüller’s lemma.

Lemma 4.6. *The sum of the interior angles of a geodesic triangle in \mathcal{U} is at most π , with equality only if the triangle is isometric to a Euclidean one. A geodesic triangle does not contain any zero of ψ in its interior.*

Lemma 4.7. *The sum of the interior angles of a geodesic quadrilateral in \mathcal{U} is at most 2π , with equality only if the quadrilateral is isometric to a Euclidean one. If the angle sum is at least π , then there is no zero of ψ in the interior of the quadrilateral and the excess angle along the sides is no more than π .*

Theorem 4.8. *Let x_0, x_1, y_1 , and y_0 be the vertices of a geodesic quadrilateral Q in \mathcal{U} . Suppose that $d(x_0, y_0) = d(x_1, y_1)$ and $\angle x_0 + \angle x_1 \geq \pi$. Then $d(y_0, y_1) \geq d(x_0, x_1)$ with equality only if Q bounds a region isometric to a Euclidean parallelogram.*

Proof. We may assume that the sides $[x_0, y_0]$ and $[x_1, y_1]$ have no singularities in their interior, for otherwise we can partition them into pairwise congruent subintervals which have no singularities in their interior and connect the partition points in pairs, thereby creating a sequence of quadrilaterals satisfying the hypotheses of the theorem, stacked on top of each other. Note that $[x_0, y_0]$ and $[x_1, y_1]$ do not intersect, since this would create a triangle with angle sum at least π at the vertices x_0 and x_1 .

The quadrilateral Q can bound at most two region. Indeed, if it bounded more, then some of them would be bigons, which is impossible by uniqueness of geodesics. If Q bounds no open set, then it is isometric to a degenerate Euclidean parallelogram, with the pairs $\{x_0, x_1\}$ and $\{y_0, y_1\}$ linked.

Suppose that Q bounds two regions, hence two triangles $\Delta_0 = x_0 z_0 y_0$ and $\Delta_1 = z_1 x_1 y_1$. The sides $[x_0, x_1]$ and $[y_0, y_1]$ possibly share an edge $[z_0, z_1]$ which we may safely disregard. We have to show that

$$d(y_0, z_0) + d(z_1, y_1) \geq d(x_0, z_0) + d(z_1, x_1).$$

Since ψ does not vanish in the interior of Q , the developing map

$$\Psi_j(\zeta) = \int_{\zeta_j}^{\zeta} \sqrt{\psi}$$

can be defined on each Δ_j . In general the developing map is only a local isometry, but in this case it turns out to be injective. The integral is only defined in the interior of Δ_j , but Ψ_j extends continuously to the boundary. The image Δ'_j of Δ_j by Ψ_j is a “fake” triangle in the plane, meaning it can have bends (or fake vertices) along its sides. Since the bends are all outwards, each side of Δ'_j is concave. Moreover, the sides of Δ'_j only intersect at actual vertices. Indeed, if two sides intersect more than once, then they create a bigon. The winding number of the triangle around any point in such a bigon is equal to -1 . However, since Ψ_j is a holomorphic function defined in a simply connected domain, it maps $\partial\Delta_j$ to a curve which winds a non-negative number of times around any point in \mathbb{C} . It follows that $\Psi_j(\partial\Delta_j)$ is a simple closed curve, which implies that Ψ_j is injective by the argument principle.

There is no loss of generality in assuming that $\Psi_j([y_j, z_j])$ is the straight line between $\Psi_j(y_j)$ and $\Psi_j(z_j)$, since this segment does not intersect the




FIGURE 4. The two triangles bounded by Q and their superposition.

other two sides and is shortest possible. Thus Δ'_j has two straight sides and a possibly bent one. Indeed, the side $\Psi_j([x_j, y_j])$ is straight since $[x_j, y_j]$ has no singularities in its interior, hence no bends. To prove the required inequality, take the fake triangles Δ'_0 and Δ'_1 and superpose them along their side of common length with a rotation or a flip so that x_0 coincides with y_1 and y_0 with x_1 , and that the angles $\angle x_0$ and $\angle x_1$ are on the same side of the superposed edge. Then the concavity of the sides and the hypothesis $\angle x_0 + \angle x_1 \geq \pi$ forces the diagonals $[y_0, z_0]$ and $[y_1, z_1]$ to intersect, say at w . Moreover, the bent sides do not cross the diagonals. This is because $\angle x_0 \geq \pi - \angle x_1 \geq \angle y_1$ and similarly $\angle x_1 \geq \angle y_0$. Now the bent side from x_0 to z_0 is strictly shorter than $d(x_0, w) + d(w, z_0)$. One way to see this is by a repeated use of the triangle inequality. If we replace a connection between consecutive bends by extending the adjacent sides until they meet, we make the bent side longer. We can repeat this until there is only one bend left, then extend the side of the bend closer to x_0 until it meets $[w, z_0]$. The result is shorter than $d(x_0, w) + d(w, z_0)$ by the triangle inequality and similarly if there is no bend at all. Thus for each $j \in \{0, 1\}$ the length of the bent side $\Psi_j([x_j, z_j])$ is bounded above by $d(x_j, w) + d(w, z_j)$. Adding these two inequalities together yields the desired result since in the superposed triangles we have $d(y_1, z_1) = d(x_0, z_1) = d(x_0, w) + d(w, z_1)$ and $d(y_0, z_0) = d(x_1, z_0) = d(x_1, w) + d(w, z_0)$. The inequality is strict in this case.

Now suppose that Q bounds a single region. The developing map Ψ for Q is an injective local isometry for the same reasons as above. Thus we can think of Q as embedded in the plane, with the two sides $[x_0, y_0]$ and $[x_1, y_1]$ straight and the other two sides possibly bent but concave. There is no loss of generality in assuming that the side joining y_0 and y_1 is straight. Suppose that the remaining side is straight as well. Then Q is just an ordinary quadrilateral in the plane. Suppose that Q is not a parallelogram. Then $\pi < \angle x_0 + \angle x_1 \leq 2\pi$ and hence $\cos(\angle x_0) + \cos(\angle x_1) < 0$. Now $d(y_0, y_1)$ is at least as large as the length of the orthogonal projection of $[y_0, y_1]$ onto the line through x_0 and x_1 . On the other hand, the length of the projection is equal to

$$d(x_0, x_1) + d(x_1, y_1) \cos(\pi - \angle x_1) - d(x_0, y_0) \cos(\angle x_0)$$

which is strictly bigger than $d(x_0, x_1)$.




FIGURE 5. The case where Q bounds a single region with bends can be reduced to the case of two triangles and the case where there are no bends.

It remains only to treat the case where the side joining x_0 and x_1 has at least one bend. In this case we have $\angle y_0 + \angle y_1 < \pi$. If there is more than one bend, then we may make the side $\Psi([x_0, x_1])$ longer by straightening consecutive bends, while making sure not to cross $[y_0, y_1]$. After we get rid of all pairs of bends we are left with two possibilities. Either there is a single bend left which is disjoint from $[y_0, y_1]$, or $\Psi([x_0, x_1])$ and $[y_0, y_1]$ share a subinterval or a point. In the latter case we can ignore the shared segment. This leaves us with two triangles, in which case we know that the strict inequality holds. Thus suppose there is only one bend, say at w . Then there is a unique straight line S which divides the sides $[x_0, y_0]$ and $[x_1, y_1]$ equally and passes through w . Then S divides Q into two triangles and one quadrilateral. The length of the bent side of Q is strictly less than the length of S by the case of two triangles, and the length of S is less than $d(y_0, y_1)$ by the case of an ordinary Euclidean quadrilateral. \square

Remark. In fact, Theorem 4.8 holds in any $\text{CAT}(0)$ space (see [4, p.184] for the definition of angles in $\text{CAT}(\kappa)$ spaces). The proof of this generalization is rather simple provided one knows the required technology. The idea is to split the quadrilateral into two triangles via a diagonal and to use comparison triangles. Since $\text{CAT}(0)$ spaces won't be used in the sequel, we preferred to stick with the above tedious but elementary proof.

We can now return to proving Theorem 4.1. We have two slit mappings $f_0, f_1 : X \rightarrow Y$ and a homotopy $H : X \times [0, 1] \rightarrow Y$ between them. The maps f_0 and f_1 extend to continuous maps from \hat{X} to \hat{Y} since they are 1-Lipschitz with respect to the complete metrics $|\varphi|^{1/2}$ and $|\psi|^{1/2}$, where $\varphi = f_0^* \psi = f_1^* \psi$. The last equality holds because $f_1 \circ f_0^{-1}$ is a local translation with respect to ψ by Theorem 3.7. We may assume that f_0 and f_1 both carry \dot{X} into \dot{Y} , since we can fill in any other punctures and treat them as points of X . Since f_0 is homotopic to f_1 on X , we must have $f_0 = f_1$ on \dot{X} . Thus we may assume that the homotopy $H : X \times [0, 1] \rightarrow Y$ extends to be constant at the punctures by Lemma 2.8.

Given $x \in X$, choose a lift of the path $t \mapsto H(x, t)$ to \mathcal{U} , let γ_x be the unique geodesic between the endpoints of this lift, let $L(x)$ be its length, and let $F(x, t) = b \circ \gamma_x(t)$. Note that $L(x)$ and $F(x, t)$ do not depend on the choice of lift, since the deck group of the branched cover $b : \mathcal{U} \rightarrow \hat{Y}$ acts by isometries on \mathcal{U} .

The maps L and F are continuous on X by Lemma 4.5 and the fact that H is continuous. We extend L and F to \hat{X} as follows. If $x \in \partial X$, then take a simply connected ball B containing x . The restriction $H : (B \cap X) \times [0, 1] \rightarrow W$ lifts to \mathcal{U} and extends continuously to $B \times \{0, 1\}$. For every $y \in B \cap X$, we have a well-defined geodesic γ_y . Let γ_x be the limit of γ_y as $y \rightarrow x$, let $L(x)$ be its length, and let $F(x, t) = b \circ \gamma_x(t)$ as before. Since H is constant on \dot{X} , we have $L(x) \rightarrow 0$ as x converges to any point in \dot{X} and we extend F by letting $F(x, t) = f_0(x)$ for every $(x, t) \in \dot{X} \times [0, 1]$. The fact that $f_1 \circ f_0^{-1}$ is a local translation allows us to analyze the local maxima of L with the aid of Theorem 4.8.

Theorem 4.9. *Suppose that $f_0 \neq f_1$ and that L has a local maximum at $x \in \hat{X}$. Then L is locally constant near x . If γ_x passes through a zero of ψ , then $x \in \partial X$ and γ_x is horizontal. If φ has a zero at x , then $x \in \partial X$ and the zero is simple.*

Proof. Suppose that $S \subset X \cup \partial X$ is a simply connected set containing x such that $f_j(S)$ is disjoint from \dot{Y} . Then there is a unique lift \tilde{f}_j of $f_j|_S$ to \mathcal{U} mapping x to $\gamma_x(j)$. We will write $\sigma_j = \tilde{f}_j(\sigma)$ for every $\sigma \in S$. Since L has a local maximum at x , there is a simply connected ball B centered at x such that $L(y) \leq L(x)$ for every $y \in B$. If $L(x) = 0$, then $L(y) = 0$ for every $y \in B$, which means that $f_0 = f_1$ on B and hence on all of X by the identity principle. The hypothesis thus implies that $L(x) > 0$. We may therefore choose the radius of B to be less than $L(x)$.

Case 0: Suppose that $x \in \dot{X}$. Then $L(x) = 0$, contradicting the above.

Case 1: Suppose that $x \in X$. Then neither $f_0(x)$ nor $f_1(x)$ is in \dot{Y} and by shrinking B we may assume that $f_0(B)$ and $f_1(B)$ are disjoint from \dot{Y} . Let $y \in B \setminus x$ and consider the quadrilateral Q formed by the vertices x_0, x_1, y_1 , and y_0 . If $\angle y_0 x_0 x_1 + \angle x_0 x_1 y_1 \geq \pi$, then

$$L(x) = d(x_0, x_1) \leq d(y_0, y_1) = L(y) \leq L(x)$$

by Theorem 4.8. Since equality holds, Q is isometric to a Euclidean parallelogram and $\angle y_0 x_0 x_1 + \angle x_0 x_1 y_1 = \pi$. Now if $\angle y_0 x_0 x_1 + \angle x_0 x_1 y_1 \leq \pi$, then we replace y by a point y' on the other side of x to form another quadrilateral Q' . More precisely, let $y' \in B \setminus x$ be such that the angle $\angle y x y'$ is at least π on either side. Since f_j is injective on B , we also have $\angle y_j x_j y'_j \geq \pi$ on either side for $j = 0, 1$. We thus get

$$\angle y'_0 x_0 x_1 \geq \pi - \angle y_0 x_0 x_1 \quad \text{and} \quad \angle x_0 x_1 y'_1 \geq \pi - \angle x_0 x_1 y_1$$

so that $\angle y'_0 x_0 x_1 + \angle x_0 x_1 y'_1 \geq 2\pi - \pi = \pi$. By the above reasoning the quadrilateral Q' with vertices x_0, x_1, y'_1 , and y'_0 is isometric to a Euclidean parallelogram. This shows that $\angle y'_0 x_0 x_1 + \angle x_0 x_1 y'_1 = \pi$ and hence $\angle y_0 x_0 x_1 + \angle x_0 x_1 y_1 = \pi$. Thus the first situation holds and $L(y) = L(x)$, i.e. L is constant on B . Suppose that Q is degenerate. This means that $[x_j, y_j] \subset \gamma_x$ for $j = 0$ or $j = 1$. Since f_j is injective, this can hold for at most one

ray $\eta \subset B$. Let $y, y' \in B \setminus \eta$ be a pair of points such that $\angle xyx' \geq \pi$ on either side. The corresponding parallelograms Q and Q' are non-degenerate and on different sides of γ_x , which implies that there is no excess angle on either side of γ_x . In other words, the interior of γ_x does not contain any zero of ψ . Suppose however that ψ has zeros at the endpoints of γ_x . Then φ has a zero at x and there exists a $y \in B \setminus x$ such that the angle between $[y_0, x_0]$ and γ_x is strictly greater than π . The corresponding quadrilateral Q is not isometric to a Euclidean parallelogram, which contradicts the above reasoning. Therefore, γ_x is completely free of singularities.

If $x \in \partial X$, then the same basic idea works. However, we distinguish a few subcases to say a bit more when φ has a zero at x .

Case 2: Suppose that $x \in \partial X$, that neither $f_0(x)$ nor $f_1(x)$ is in \dot{Y} , and that φ does not have a zero at x . Then let y, y' be on opposite sides of x in the horizontal segment $B \cap \partial X$. By the same argument as in Case 1, the corresponding quadrilaterals Q and Q' are parallelograms. If Q and Q' are degenerate, then for one of $j = 0$ or $j = 1$ we have $[x_j, y_j] \subset \gamma_x$ and $[x_{j+1}, y'_{j+1}] \subset \gamma_x$ where the index is taken mod 2. Moreover, $[x_j, y_j]$ and $[x_{j+1}, y'_{j+1}]$ are horizontal since f_j is a local translation. It follows that for every $y'' \in B$, the angles $\angle y''_0 x_0 x_1$ and $\angle x_0 x_1 y''_1$ are complementary so that $L(y'') = L(x)$. If Q or Q' is non-degenerate, then neither is the other one, and the interior of γ_x is free of zeros. Thus the slope of γ_x is constant. This implies that $\angle y''_0 x_0 x_1 + \angle x_0 x_1 y''_1 = \pi$ for every $y'' \in B$ so that $L(y'') = L(x)$. In either case L is constant on B .

Case 3: Suppose that $x \in \partial X$, that neither $f_0(x)$ nor $f_1(x)$ is in \dot{Y} , and that φ has a zero at x . Then the total angle at x is at least 2π . Thus for every $y \in B \setminus x$, there is a $y' \in B \setminus x$ such that $\angle xyx' \geq \pi$ on either side. Then we can argue in exactly the same way as in Case 1 and we conclude that L is constant on B . Suppose that the quadrilateral Q with vertices x_0, x_1, y_1 , and y_0 is degenerate for some $y \in B$. Then $[x_j, y_j] \subset \gamma_x$ for $j = 0$ or $j = 1$. Since f_j is injective on $B \cap X$, this can hold for y in at most one ray $\eta \subset B \cap X$. Thus we can still find two points $y, y' \in B \cap X \setminus \eta$ such that $\angle xyx' \geq \pi$ on either side. The corresponding parallelograms Q and Q' are non-degenerate and on different sides of γ_x , so that the interior of γ_x is free of singularities. Suppose that there is some $y \in B \setminus x$ such that $\angle y_0 x_0 x_1 > \pi$ on either side. Then $L(y) > L(x)$ by Theorem 4.8, which is a contradiction. It follows that the zero of φ at x is simple (i.e. half of a double zero) and that the initial portion of γ_x is the horizontal trajectory bisecting $\tilde{f}_0(B \cap X)$.

Case 4: Suppose that $x \in \partial X$ and that $f_j(x) \in \dot{Y}$ for some $j \in \{0, 1\}$. Since $B \cap X$ is simply connected, the restriction $f_j|_{B \cap X}$ lifts to a map $\tilde{f}_j : B \cap X \rightarrow \mathcal{U}$. Moreover the lift extends continuously to $B \cap \partial X$ since the corresponding part of the frontier of $\tilde{f}_j(B \cap X)$ is horizontal, hence a piecewise analytic curve. Thus one of Case 2 or Case 3 applies without change. \square

We conclude that L is constant.

Corollary 4.10. *The function L is constant.*

Proof. Let M be the closed subset of \widehat{X} where L attains its maximum value. Every point of M is a local maximum of L . By Theorem 4.9, L is locally constant near every local maximum, so that M is open. Since \widehat{X} is connected, $M = \widehat{X}$. \square

Hence every point of \widehat{X} is a local maximum of L , so that the conclusions of Theorem 4.9 hold everywhere.

Corollary 4.11. *Suppose that $f_0 \neq f_1$. Then \dot{X} is empty and φ has no zeros in X and at most simple zeros in ∂X . Moreover, γ_x is horizontal for every $x \in \widehat{X}$.*

Proof. Since L is constant, every $x \in \widehat{X}$ is a local maximum of L . If \dot{X} is non-empty, then $f_0 = f_1$. By Theorem 4.9, if φ has a zero at $x \in X \cup \partial X$, then $x \in \partial X$ and the zero is simple.

For every $x \in X$, the geodesic γ_x is free of singularities. Moreover, for every y near x , the geodesic γ_y forms a parallelogram with γ_x . Thus the slope of γ_x is locally constant as a function of x , hence constant on X . Since the map $x \mapsto \gamma_x$ is locally continuous on $X \cup \partial X$, the slope is constant on all $X \cup \partial X$.

Let $y \in \widehat{Y}$ be a zero of ψ or a point of \dot{Y} where ψ does not have a pole. Since the map $f_0 : \widehat{X} \rightarrow \widehat{Y}$ is surjective, there is some $x \in \widehat{X}$ with $f_0(x) = y$. The corresponding geodesic γ_x has a zero at its starting point. By Theorem 4.9, $x \in \partial X$ and γ_x is horizontal. Thus all other geodesics are horizontal.

If ψ has no zeros in \widehat{Y} , then \widehat{Y} is either an annulus or a torus. Suppose that \widehat{Y} is an annulus. If the geodesics γ_x are not horizontal, then they point away from Y on one of the two boundary components, which is of course impossible. If \widehat{Y} is a torus, then \dot{Y} is non-empty since Y is hyperbolic. If there is a point of \dot{Y} which is not a pole of ψ , then all the geodesics γ_x are horizontal by the previous paragraph. On the other hand, if ψ has a pole in \dot{Y} then it has at least one zero in \widehat{Y} and once again all geodesics are horizontal. \square

We deduce that F is the desired isotopy from f_0 to f_1 .

Theorem 4.12. *If $f_0 \neq f_1$, then the map $F : X \times [0, 1] \rightarrow Y$ is an isotopy from f_0 to f_1 through slit mappings.*

Proof. For every $x \in X$, the geodesic γ_x is free of zeros and nearby geodesics are parallel of the same length, which shows that $f_t(x) = F(x, t)$ is a local translation. We have to show that f_t is injective and that its image is contained in Y .

Since f_t is a local isometry, the area of its image is at most the area of X , with equality only if f is injective. Thus to show that f is injective it suffices

to show that it is almost surjective. Recall that there is a full measure subset $U \subset X$ such that through every point of U passes a bi-infinite horizontal trajectory of φ . For every bi-infinite geodesic $\alpha \subset U$, we have $f_t(\alpha) = f_0(\alpha)$ since $f_0(\alpha)$ is horizontal for ψ and F moves points horizontally. It follows that $f_t(U) = f_0(U)$. But $f_0(U)$ has full measure in Y since U has full measure in X and $f_0(X)$ has full measure in Y . Therefore $f_t(X)$ has full measure in \widehat{Y} and hence f_t is injective. Since f_t is a local translation and $f_t(X)$ is dense in \widehat{Y} , it follows that f_t is a slit mapping with respect to ψ by Lemma 3.6.

It remains to show that $f_t(X)$ is disjoint from \dot{Y} . Suppose that $f_t(x) \in \dot{Y}$ for some $x \in X$. If $f_t(x)$ is not a pole of ψ in \widehat{Y} , then $\gamma_x(t)$ is a zero of ψ in \mathcal{U} . By Theorem 4.9, this implies that $x \in \partial X$, contradiction. Thus $f_t(x)$ is a simple pole of ψ . Since f_t is a local translation, φ has a simple pole at x . This is also impossible since $\varphi = f_0^* \psi$ is holomorphic in X . Thus the open set $f_t(X)$ is contained in $Y \cup \partial Y$, and hence in Y . \square

This completes the proof of Theorem 4.1. We record the following consequence for later reference.

Lemma 4.13. *Suppose that $\text{CEmb}(X, Y, h)$ contains a slit mapping. Then for every $x \in X$, the evaluation map which sends $f \in \text{CEmb}(X, Y, h)$ to $f(x) \in Y$ is a homeomorphism onto its image $V(x)$ which is either a closed regular horizontal trajectory, a compact horizontal arc, or a point.*

Proof. The evaluation map is continuous by definition of the compact-open topology. We now show that it is injective. Suppose that $f_0(x) = f_1(x)$ for distinct $f_0, f_1 \in \text{CEmb}(X, Y, h)$ and some $x \in X$. Let $F : X \times [0, 1] \rightarrow Y$ be the isotopy from f_0 to f_1 constructed in the proof of Theorem 4.1. Then $t \mapsto \alpha(t) = F(x, t)$ is a closed regular horizontal trajectory for ψ . Let A be an annular neighborhood of α foliated by closed trajectories which is symmetric about α . Since $f_0(x)$ is a fixed point of the local translation $f_1 \circ f_0^{-1}$, we have that $f_1 \circ f_0^{-1}$ is either equal to the identity near $f_0(x)$ or a half-turn around $f_0(x)$. Suppose it is locally a half-turn. Let σ be the conformal involution of A which fixes $f_0(x)$ and permutes the two boundary components. By the identity principle $f_1 \circ f_0^{-1} = \sigma$ on the connected component C of $A \cap f_0(X)$ containing f_0 . Let β' be any simple close curve contained in C and winding once around A and let $\beta = f_0^{-1}(\beta')$. Then $f_1(\beta)$ and $f_0(\beta)$ are not homotopic in Y , for they are inverses of one another. This is a contradiction. Otherwise, $f_1 \circ f_0^{-1}$ is equal to the identity near $f_0(x)$ and hence on all of $f_0(X)$ by the identity principle. This contradicts the hypothesis that $f_0 \neq f_1$ and we conclude that the evaluation map is injective. Any injective continuous map from a compact space to a Hausdorff space is a homeomorphism onto its image.

Since the space $\text{CEmb}(X, Y, h)$ is compact, its image $V(x)$ by the evaluation map is compact. Moreover $V(x)$ is path-connected and contained in a

regular horizontal trajectory by Theorem 4.1. Thus $V(x)$ is either a closed trajectory, a segment, or a point. \square

5. THE MODULUS OF EXTENSION

Let $h : X \rightarrow Y$ be a non-trivial and non-parabolic embedding. Recall that the space $\text{CEmb}(X, Y, h)$ is the set of all conformal embeddings homotopic to h equipped with the compact-open topology. This space is compact by Lemma 3.2. If X has empty ideal boundary, then $\text{CEmb}(X, Y, h)$ is either empty or a singleton.

Lemma 5.1. *If ∂X is empty, then $\text{CEmb}(X, Y, h)$ has at most one element.*

Proof. Every $f \in \text{CEmb}(X, Y, h)$ extends to a conformal isomorphism \hat{f} between the closed surfaces \hat{X} and \hat{Y} by Riemann's removable singularity theorem. If f and g are homotopic, then so are \hat{f} and \hat{g} . The composition $\hat{g}^{-1} \circ \hat{f}$ is thus a conformal automorphism of \hat{X} homotopic to the identity. If \hat{X} has genus at least 2, then it is hyperbolic and $\hat{f} = \hat{g}$ by Lemma 3.8. If \hat{X} has genus 0, then \hat{X} must contain at least 3 points for X to be hyperbolic. Any conformal automorphism of the Riemann sphere fixing 3 points is the identity. If \hat{X} has genus 1, then \hat{X} has at least 1 point. Any conformal automorphism of a torus homotopic to the identity is a translation, and a translation with a fixed point is the identity. \square

In light of this result, **we assume that ∂X and $\text{CEmb}(X, Y, h)$ are non-empty from now on.** For each component C of ∂X , choose an analytic parametrization $\zeta_C : S^1 \rightarrow C$. For every $r \in (0, \infty]$ and every component C of ∂X , glue a copy of the cylinder $S^1 \times [0, r)$ to $X \cup \partial X$ along $S^1 \times \{0\}$ with the map ζ_C (see Figure 6). We denote the resulting surface by X_r . We also let $X_0 = X$. If $\rho \leq r$, then the inclusion $[0, \rho) \subset [0, r)$ induces a conformal inclusion $X_\rho \subset X_r$. Note that for every $r \in [0, \infty]$, there is a homeomorphism $H : X_r \rightarrow X$ which when followed by the inclusion $X \subset X_r$ is homotopic to the identity. We will abuse notation and write $\text{CEmb}(X_r, Y, h)$ instead of $\text{CEmb}(X_r, Y, h \circ H)$.




FIGURE 6. The surface X_r is obtained by gluing a cylinder of modulus r to each ideal boundary component of X .

Given $f \in \text{CEmb}(X, Y, h)$, we define the *modulus of extension* of f as $\mathfrak{m}(f) = \sup\{r \in [0, \infty] \mid f \text{ extends to a conformal embedding } X_r \hookrightarrow Y\}$.

It is easy to see that the supremum is realized and that the extension is unique.

Lemma 5.2. *For every $f \in \text{CEmb}(X, Y, h)$, there is a unique conformal embedding $\hat{f} : X_{\mathbf{m}(f)} \rightarrow Y$ extending f .*

Proof. Let $r = \mathbf{m}(f)$, let $r_n \nearrow r$, and let $f_n : X_{r_n} \rightarrow Y$ be a conformal embedding extending f . Every $x \in X_r$ is contained in X_{r_n} when n is large enough since $X_{r_n} \nearrow X_r$. Define $\hat{f}(x) = f_n(x)$. This does not depend on n since $f_j = f_{j+1}$ on X_{r_j} by the identity principle. The function \hat{f} is holomorphic and injective on X_r because each f_n is holomorphic and injective. The uniqueness of \hat{f} follows from the identity principle. \square

Similarly, \mathbf{m} is upper semi-continuous on $\text{CEmb}(X, Y, h)$.

Lemma 5.3. *The modulus of extension \mathbf{m} is upper semi-continuous.*

Proof. Suppose that $f_n \rightarrow f$ in $\text{CEmb}(X, Y, h)$. Pass to a subsequence so that $r_n = \mathbf{m}(f_n)$ converges to some $r \in [0, \infty]$. We have to show that $\mathbf{m}(f) \geq r$. If $r = 0$, then there is nothing to prove, so we may assume that $r > 0$. Let $\rho < r$ and let \hat{f}_n be the maximal injective holomorphic extension of f_n . If n is large enough, then $\rho \leq r_n$, and we let g_n be the restriction of \hat{f}_n to X_ρ . Since g_n belongs to the compact space $\text{CEmb}(X_\rho, Y, h)$, we may pass to a subsequence such that $g_n \rightarrow g$ for some $g \in \text{CEmb}(X_\rho, Y, h)$. Now the restriction of g_n to X is equal to f_n , which by hypothesis converges to f . Thus the restriction of g to X is equal to f . In other words, g is a conformal embedding extending f , so that $\mathbf{m}(f) \geq \rho$. Since $\rho < r$ was arbitrary we have $\mathbf{m}(f) \geq r$. \square

Since \mathbf{m} is upper semi-continuous and $\text{CEmb}(X, Y, h)$ is compact, \mathbf{m} achieves its maximum at some $f \in \text{CEmb}(X, Y, h)$. We will see that any such maximum is the restriction of a slit mapping. We first need to show that any limit of Teichmüller embeddings is itself a Teichmüller embedding.

Definition 5.4. Let $X_n \in \mathcal{T}^\#(X)$ and $Y_n \in \mathcal{T}^\#(Y)$ be such that $X_n \rightarrow X$ and $Y_n \rightarrow Y$ as $n \rightarrow \infty$, and let $\sigma_n : X_n \rightarrow X$ and $\tau_n : Y_n \rightarrow Y$ be the Teichmüller maps that respect the markings. Let $h_n : X_n \rightarrow Y_n$ and $h : X \rightarrow Y$ be any maps. We say that $h_n \rightarrow h$ if $\tau_n \circ h_n \circ \sigma_n^{-1} \rightarrow h$ locally uniformly on X .

Lemma 5.5. *Let $h : X \rightarrow Y$ be an embedding between finite Riemann surfaces. Suppose that $X_n \rightarrow X$ in $\mathcal{T}^\#(X)$ and $Y_n \rightarrow Y$ in $\mathcal{T}^\#(Y)$, and let $f_n : X_n \rightarrow Y_n$ be a sequence of Teichmüller embeddings homotopic to h . Then there exists a Teichmüller embedding $f : X \rightarrow Y$ homotopic to h such that $f_n \rightarrow f$ after passing to a subsequence.*

Proof. After passing to a subsequence, we may assume that $\text{Dil}(f_n) \rightarrow K$ for some $K \geq 1$. By Lemma 3.2, we may also assume that $f_n \rightarrow f$ for some K -quasiconformal embedding $f : X \rightarrow Y$ homotopic to h . We claim that

f has minimal dilatation in its homotopy class. If not, let $g : X \rightarrow Y$ be a quasiconformal embedding homotopic to f with $\text{Dil}(g) < \text{Dil}(f)$. Then let $\sigma_n : X_n \rightarrow X$ and $\tau_n : Y_n \rightarrow Y$ be the Teichmüller maps that respect the markings. By hypothesis we have $\text{Dil}(\sigma_n) \rightarrow 1$ and $\text{Dil}(\tau_n) \rightarrow 1$. Thus if $\varepsilon > 0$ is small enough and n is large enough, then

$$\text{Dil}(\tau_n^{-1} \circ g \circ \sigma_n) \leq (1 + \varepsilon) \text{Dil}(g) < \text{Dil}(f) \leq K.$$

It follows that $\text{Dil}(\tau_n^{-1} \circ g \circ \sigma_n) < \text{Dil}(f_n)$ if n is large enough, which contradicts Theorem 3.7. Thus f has minimal dilatation in its homotopy class. If f is not conformal, then it is a Teichmüller embedding by Theorem 3.5, and we are done.

Suppose that f is conformal. Let ψ_n be the terminal quadratic differential of f_n , and let $g_n : Y_n \rightarrow Y'_n$ be the Teichmüller map of dilatation e^2 and initial quadratic differential ψ_n . Since $d(Y_n, Y'_n) = 1$ for every n and $Y_n \rightarrow Y$, and since closed balls in $\mathcal{T}^\#(Y)$ are compact, we may pass to a subsequence such that $Y'_n \rightarrow Y'$ for some $Y' \in \mathcal{T}^\#(Y)$ with $d(Y, Y') = 1$. Let $g : Y \rightarrow Y'$ be the Teichmüller map that respects the marking. By a standard argument similar to the one above, $g_n \rightarrow g$. Indeed, every subsequence of $\{g_n\}$ has a converging subsequence. Any limit of any subsequence is a quasiconformal homeomorphism with minimal dilatation in its homotopy class, hence equal to g .

By construction, $g_n \circ f_n : X_n \rightarrow Y'_n$ is a Teichmüller embedding. Moreover, $g_n \circ f_n \rightarrow g \circ f$. As in the first paragraph of this proof, $g \circ f$ has minimal dilatation in its homotopy class. This dilatation is equal to $e^2 > 1$, so that $g \circ f$ is a Teichmüller embedding. Since f is conformal, the terminal quadratic differential of $g \circ f$ is equal to the terminal quadratic differential of g . It follows that $f = g^{-1} \circ (g \circ f)$ is a slit mapping with respect to the initial quadratic differential of g . \square

We obtain the following characterization of the global maxima of \mathfrak{m} .

Theorem 5.6. *Let f be a global maximum of \mathfrak{m} with $\mathfrak{m}(f) < \infty$ and let \hat{f} be the extension of f to $X_{\mathfrak{m}(f)}$. Then \hat{f} is a slit mapping.*

Proof. Let $R = \mathfrak{m}(f)$. For every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y$ whose restriction to X is homotopic to h , for otherwise we would have $\mathfrak{m}(g|_X) \geq r > R = \mathfrak{m}(f)$. By Ioffe's theorem, there exists a Teichmüller embedding $g_r : X_r \rightarrow Y$ whose restriction to X is homotopic to h . It is easy to see that X_r moves continuously in $\mathcal{T}^\#(X)$ as a function of $r \in [0, \infty)$. By Lemma 5.5, we can extract a limiting Teichmüller embedding $g : X_R \rightarrow Y$ from some subsequence of g_r as $r \rightarrow R$. Since \hat{f} is conformal and homotopic to g , Ioffe's theorem implies that \hat{f} is itself a Teichmüller embedding. A conformal Teichmüller embedding is a slit mapping. \square

Observe that every end of the surface X_∞ is a puncture since the cylinder $S^1 \times [0, \infty)$ is conformally isomorphic to $\mathbb{D} \setminus 0$. Thus if $\mathfrak{m}(f) = \infty$, then the extension \hat{f} extends further to a conformal isomorphism between $\widehat{X_\infty}$ and

\widehat{Y} . In particular, $Y \setminus \widehat{f}(X_\infty)$ is finite, so that \widehat{f} is a slit mapping with respect to any $\psi \in \mathcal{Q}^+(Y)$. This is unless Y is the triply punctured sphere in which case $\mathcal{Q}^+(Y)$ is empty. Thus the hypothesis $\mathfrak{m}(f) < \infty$ in the above theorem is superfluous unless Y is the triply punctured sphere. In the latter case \widehat{f} is nevertheless unique in its homotopy class by Lemma 5.1.

We can in fact strengthen Theorem 5.6 by replacing the word “global” with “local”. The proof requires another lemma similar to Lemma 5.5.

Lemma 5.7. *Suppose that $Y_n \nearrow Y$, where the inclusion $Y_n \subset Y$ is homotopic to a homeomorphism, and let $f_n : X \rightarrow Y_n$ be a sequence of Teichmüller embeddings. Then there exists a Teichmüller embedding $f : X \rightarrow Y$ such that $f_n \rightarrow f$ after passing to a subsequence.*

Proof. Let K_n be the dilatation of f_n . The sequence K_n is non-increasing and thus converges to a limit $K \geq 1$. After passing to a subsequence, f_n converges to a K -quasiconformal embedding $f : X \rightarrow Y$.

We claim that f has minimal dilatation in its homotopy class. If not, let $g : X \rightarrow Y$ be a quasiconformal embedding homotopic to f with $\text{Dil}(g) < \text{Dil}(f)$. If $r > 0$ is small enough, then there is a quasiconformal homeomorphism $\sigma : X_r \rightarrow X$ of dilatation strictly smaller than $\text{Dil}(f)/\text{Dil}(g)$. Thus the quasiconformal embedding f' consisting of the inclusion $X \hookrightarrow X_r$ followed by $g \circ \sigma$ has dilatation strictly less than $\text{Dil}(f)$. Then $f'(X) \subset Y_n$ whenever n is large enough. Indeed, the ends of $f'(X)$ of finite modulus are compactly contained in Y since they are surrounded by the union of collars $g \circ \sigma(X_r \setminus X)$. Thus the only way to go to infinity in the closure $\overline{f'(X)}$ relative to Y is via punctures of $f'(X)$ that are also punctures of Y . For every such puncture p and every $n \in \mathbb{N}$, a neighborhood of p in Y is contained in Y_n . This is because f_n is quasiconformal and as such it cannot map punctures of X to ends of Y_n with finite modulus. Since $Y_n \nearrow Y$, the inclusion $f'(X) \subset Y_n$ holds for n large enough. But the inequality $\text{Dil}(f') < \text{Dil}(f) \leq \text{Dil}(f_n)$ contradicts the hypothesis that f_n is a Teichmüller embedding and hence of minimal dilatation in its homotopy class.

If f is not conformal, then it is a Teichmüller embedding by Theorem 3.5, and we are done. Thus suppose that f is conformal but is not a slit mapping. By Theorem 5.6, there exists an $r > 0$ and a conformal embedding $g : X_r \rightarrow Y$ whose restriction to X is homotopic to f . By the above argument, we have $g(X_{r/2}) \subset Y_n$ if n is large enough, and thus $g(X) \subset Y_n$ with complement having non-empty interior. On the other hand, the restriction $g|_X : X \rightarrow Y_n$ is conformal and homotopic to the Teichmüller embedding $f_n : X \rightarrow Y_n$. By Theorem 3.7, $g|_X$ is a slit mapping so that $Y_n \setminus g(X)$ has empty interior, contradiction. \square

We come to the main result of this section, which is that every local maximum of \mathfrak{m} is the restriction of a slit mapping.

Theorem 5.8. *Let f be a local maximum of \mathfrak{m} with $\mathfrak{m}(f) < \infty$, and let \widehat{f} be the extension of f to $X_{\mathfrak{m}(f)}$. Then \widehat{f} is a slit mapping.*

Proof. Let $R = \mathfrak{m}(f)$. We first show that the complement $Y \setminus \widehat{f}(X_R)$ is horizontal for some meromorphic quadratic differential on Y , and is in particular an analytic graph. Let $\{x_1, x_2, \dots\}$ be a dense subset of X .

Claim. *There exists an $n \in \mathbb{N}$ such that if $g \in \text{CEmb}(X, Y, h)$ satisfies $g(x_j) = f(x_j)$ for every $j \in \{1, \dots, n\}$, then $\mathfrak{m}(g) \leq \mathfrak{m}(f)$.*

Proof of Claim. Suppose on the contrary that for every $n \in \mathbb{N}$ there exists an element g_n of $\text{CEmb}(X, Y, h)$ satisfying $g_n(x_j) = f(x_j)$ for every $j \in \{1, \dots, n\}$ such that $\mathfrak{m}(g_n) > \mathfrak{m}(f)$. As $\text{CEmb}(X, Y, h)$ is compact, every subsequence of $\{g_n\}_{n=1}^\infty$ has a subsequence converging to some $g \in \text{CEmb}(X, Y, h)$. Any limit g agrees with f on the dense set $\{x_1, x_2, \dots\}$, and hence is equal to f . Thus $g_n \rightarrow f$ with $\mathfrak{m}(g_n) > \mathfrak{m}(f)$. This contradicts the hypothesis that f is a local maximum of \mathfrak{m} . \square

Let n be as in the claim, and let $P = \{x_1, x_2, \dots, x_n\}$. Then for every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y$ homotopic to f rel P . By Ioffe's theorem, there exists a Teichmüller embedding $g_r : X_r \setminus P \rightarrow Y \setminus f(P)$ homotopic to f rel P .

Let g be any limit of any subsequence of g_r as $r \searrow R$. Then $g : X_R \setminus P \rightarrow Y \setminus f(P)$ is a Teichmüller embedding by Lemma 5.5. Since \widehat{f} is conformal and homotopic to g rel P , Ioffe's theorem implies that \widehat{f} is itself a slit mapping, considered as a map from $X_R \setminus P$ to $Y \setminus f(P)$. Therefore the complement $Y \setminus \widehat{f}(X_R)$ is a finite union of horizontal arcs for some meromorphic quadratic differential on Y , possibly with simple poles on the set $f(P)$.

Let $\Gamma = Y \setminus \widehat{f}(X_R)$, let $\{y_1, y_2, \dots\}$ be a dense subset of the graph Γ minus its vertices, and fix a Riemannian metric on Y , say the hyperbolic one.

Claim. *There exists a $k \in \mathbb{N}$ such that for every $r > R$ and every $\varepsilon > 0$, there is no conformal embedding $g : X_r \rightarrow Y$ homotopic to h whose image is disjoint from the balls $\overline{B}(y_1, \varepsilon), \dots, \overline{B}_\varepsilon(y_k, \varepsilon)$.*

Proof of Claim. Suppose that for every $k \in \mathbb{N}$ there exist an $r_k > R$, an $\varepsilon_k > 0$, and a conformal embedding $g_k : X_{r_k} \rightarrow Y$ whose restriction to X is homotopic to h such that $g_k(X_{r_k})$ is disjoint from the balls $\overline{B}(y_1, \varepsilon_k), \dots, \overline{B}(y_k, \varepsilon_k)$. We may assume that $r_k \rightarrow R$ and $\varepsilon_k \rightarrow 0$. Let g be any limit of any subsequence of the sequence $\{g_k\}$. Then $g(X_R)$ is disjoint from the set $\{y_1, y_2, \dots\}$ and hence from its closure Γ , so that $\widehat{f}^{-1} \circ g : X_R \rightarrow X_R$ is a conformal embedding homotopic to the identity. If X_R is not an annulus, then Lemma 3.8 implies that $g = \widehat{f}$ and hence $g_k \rightarrow \widehat{f}$. If X_R is an annulus, then we may pre-compose each g_k by a rotation so that we still get $g_k \rightarrow \widehat{f}$. Since $\mathfrak{m}(g_k|_X) \geq r_k > R = \mathfrak{m}(f)$, this contradicts the hypothesis that f is a local maximum of \mathfrak{m} . \square

Let k be as in the last claim, and let $Q = \{y_1, \dots, y_k\}$. For each $\varepsilon > 0$, let $Y_\varepsilon = Y \setminus (\overline{B}(y_1, \varepsilon) \cup \dots \cup \overline{B}(y_k, \varepsilon))$. Let $\iota_\varepsilon : Y_\varepsilon \rightarrow Y \setminus Q$ be a homeomorphism homotopic to the inclusion map, and let $h_\varepsilon = \iota_\varepsilon^{-1} \circ \hat{f}$. The embedding $h_\varepsilon : X_R \rightarrow Y_\varepsilon$ followed by the inclusion $Y_\varepsilon \hookrightarrow Y$ is homotopic to h . By the claim, for every $r > R$, there is no conformal embedding $g : X_r \rightarrow Y_\varepsilon$ homotopic to h_ε . Therefore, there is a Teichmüller embedding $g'_\varepsilon : X_r \rightarrow Y_\varepsilon$ homotopic to h_ε . Letting $r \rightarrow R$, we can extract a limiting Teichmüller embedding $g_\varepsilon : X_R \rightarrow Y_\varepsilon$ by Lemma 5.5.

Since $Y_\varepsilon \nearrow (Y \setminus Q)$ as $\varepsilon \searrow 0$, we can apply Lemma 5.7 and obtain a Teichmüller embedding $g : X_R \rightarrow Y \setminus Q$ as a limit of a subsequence of $\{g_\varepsilon\}$. Since $\hat{f} : X_R \rightarrow Y \setminus Q$ is homotopic to g , it is a slit mapping with respect to some $\psi \in \mathcal{Q}^+(Y \setminus Q)$. Thus ψ is meromorphic on Y with at most simple poles on Q . Moreover, the graph $\Gamma = Y \setminus \hat{f}(X_R)$ is horizontal for ψ . Since every point of Q is contained in the interior of an edge of Γ , the quadratic differential ψ cannot have simple poles on Q . Indeed, there is only one horizontal ray emanating from any simple pole. Therefore, ψ is holomorphic on Y and $\hat{f} : X_R \rightarrow Y$ is an honest slit mapping. \square

In particular, every conformal embedding which is not a slit mapping can be approximated by a sequence of conformal embeddings, each extending by some amount.

Corollary 5.9. *Let $g \in \text{CEmb}(X, Y, h)$. Then there exists a sequence g_n converging to g in $\text{CEmb}(X, Y, h)$ such that g_n extends to a conformal embedding of X_{r_n} into Y for some $r_n > 0$ unless g is a slit mapping.*

Proof. If $\mathfrak{m}(g) > 0$, then we can take $g_n = g$. If $\mathfrak{m}(g) = 0$ but g is not a local maximum of \mathfrak{m} , then there exists a sequence $g_n \rightarrow g$ with $\mathfrak{m}(g_n) > 0$. If $\mathfrak{m}(g) = 0$ and g is a local maximum of \mathfrak{m} , then g is a slit mapping by the previous theorem. \square

A strong converse to Theorem 5.8 holds due to Ioffe's theorem.

Lemma 5.10. *Suppose that $g : X_r \rightarrow Y$ is a slit mapping such that $g|_X$ is homotopic to h . Then $g|_X$ is a global maximum of \mathfrak{m} .*

Proof. First observe that $\mathfrak{m}(g|_X) \geq r$. Suppose that $\mathfrak{m}(f) \geq \mathfrak{m}(g|_X)$ for some element f of $\text{CEmb}(X, Y, h)$ and let \hat{f} be the maximal extension of f . Then $\hat{f}|_{X_r}$ is homotopic to g . By Ioffe's theorem, $\hat{f}|_{X_r}$ is a slit mapping. In particular, the complement of $\hat{f}(X_r)$ has empty interior in Y . Therefore $X_{\mathfrak{m}(f)} \setminus X_r$ is empty so that $\mathfrak{m}(f) \leq r \leq \mathfrak{m}(g|_X)$. \square

Furthermore, the almost rigidity of slit mappings implies that the set of local maxima of \mathfrak{m} is path-connected.

Lemma 5.11. *The set M of all local maxima of \mathfrak{m} is homeomorphic to either a point, a compact interval, or a circle, and \mathfrak{m} is constant on M .*

Proof. Suppose first that there is some $f \in M$ such that $\mathbf{m}(f) < \infty$. Then by Theorem 5.8, the maximal extension \hat{f} is a slit mapping. By Lemma 5.10, f is a global maximum of \mathbf{m} . In particular, $\mathbf{m}(g) < \infty$ for every $g \in M$ and thus every $g \in M$ is a global maximum of \mathbf{m} . In particular, \mathbf{m} is constant on M , say equal to R . The map $M \rightarrow \text{CEmb}(X_R, Y, h)$ defined by $f \mapsto \hat{f}$ is a homeomorphism with inverse $g \mapsto g|_X$. By Lemma 4.13, the evaluation map $\text{ev}_x : \text{CEmb}(X_R, Y, h) \rightarrow Y$ is a homeomorphism onto its image for every $x \in X_R$, and its image is either a point, a compact interval, or a circle.

Otherwise, \mathbf{m} is constant equal to ∞ on M . In this case M is homeomorphic to $\text{CEmb}(X_\infty, Y, h)$, which is a point by Lemma 5.1. \square

The fact that \mathbf{m} has a connected plateau of local maxima easily implies that the space $\text{CEmb}(X, Y, h)$ is connected.

Theorem 5.12. *The space $\text{CEmb}(X, Y, h)$ is connected.*

Proof. Suppose that $\text{CEmb}(X, Y, h) = E_0 \cup E_1$ where E_0 and E_1 are disjoint non-empty closed sets. Then each of E_0 and E_1 is both compact and open. Since \mathbf{m} is upper semi-continuous, the restriction $\mathbf{m}|_{E_j}$ attains its maximum at some $f_j \in E_j$. Then f_j is a local maximum of \mathbf{m} since E_j is open. By Lemma 5.11, f_0 and f_1 are both contained in a connected subset M of $\text{CEmb}(X, Y, h)$. On the other hand, $M = (M \cap E_0) \cup (M \cap E_1)$ is disconnected, which is a contradiction. \square

6. THE BLOB AND ITS BOUNDARY

For Sections 6 to 10, **we assume that the embedding $h : X \rightarrow Y$ is generic**. By definition, this means that the induced homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ has non-abelian image. Let $\text{Map}(X, Y, h)$ be the space of all continuous maps $f : X \rightarrow Y$ homotopic to h . The following lemma shows that for every $x \in X$ and every $f \in \text{Map}(X, Y, h)$, there is a well-defined way to lift the image point $f(x)$ to the universal covering space of Y .

Lemma 6.1. *Let $f \in \text{Map}(X, Y, h)$ and let $H : X \times [0, 1] \rightarrow Y$ be a homotopy from h to f . Then for every $x \in X$, the homotopy class rel endpoints of the path $t \mapsto H(x, t)$ does not depend on the choice of H .*

Proof. We use two standard facts about hyperbolic surfaces:

- every abelian subgroup of $\pi_1(Y)$ is cyclic;
- every non-trivial element in $\pi_1(Y)$ is the positive power of a unique primitive element.

Let G be any homotopy from h to f . By composing H with G ran backwards, we get a homotopy from h to itself, hence a map $F : X \times S^1 \rightarrow Y$. Suppose that the loop $\gamma(t) = F(x, t)$ is not trivial in $\pi_1(Y, h(x))$. Then it is equal to β^k for some primitive element β and some $k > 0$.

Let α be any loop in X based at x . Then the map $S^1 \times S^1 \rightarrow Y$ given by $(s, t) \mapsto F(\alpha(s), t)$ induces a homomorphism of \mathbb{Z}^2 into $\pi_1(Y, h(x))$. The

image of this homomorphism is cyclic, and contains both $[h \circ \alpha]$ and $[\gamma]$. From the existence and uniqueness of primitive roots in $\pi_1(Y, h(x))$, it follows that $[h \circ \alpha] = \beta^j$ for some $j \in \mathbb{Z}$.

Since α is arbitrary, we deduce that the image of the homomorphism $\pi_1(h) : \pi_1(X, x) \rightarrow \pi_1(Y, h(x))$ is contained in the cyclic group $\langle \beta \rangle$. This contradicts the hypothesis that h is generic. We conclude that the loop $\gamma(t) = F(x, t)$ is null-homotopic. Equivalently, the paths $t \mapsto H(x, t)$ and $t \mapsto G(x, t)$ are homotopic rel endpoints. \square

Fix some point $x \in X$ and a universal covering map $\pi_Y : \mathbb{D} \rightarrow Y$, and let $b \in \mathbb{D}$ be such that $\pi_Y(b) = h(x)$.

Definition 6.2. Let $f \in \text{Map}(X, Y, h)$. We define $\text{lift}_x(f)$ as the endpoint of the unique lift of the path $t \mapsto H(x, t)$ based at b , where H is an arbitrary homotopy from h to f . This does not depend on H by Lemma 6.1.

Lemma 6.3. *The map $\text{lift}_x : \text{Map}(X, Y, h) \rightarrow \mathbb{D}$ is continuous.*

Proof. Let $f_n, f \in \text{Map}(X, Y, h)$ be such that $f_n \rightarrow f$. Let $K \subset X$ be a compact deformation retract of X containing x . Let $\varepsilon > 0$ be smaller than the injectivity radius of Y on $f(K)$, and let n be large enough so that $|f_n - f| < \varepsilon$ on K . For every $(\xi, t) \in K \times [0, 1]$, let $G'(\xi, t)$ be the point at proportion t along the unique shortest length geodesic between $f(\xi)$ and $f_n(\xi)$. This gives a continuous homotopy from $f|_K$ to $f_n|_K$. By composing the deformation retraction $X \rightarrow K$ with G' , we get a homotopy G from f to f_n moving points of K by distance at most ε .

Given any homotopy H from h to f , the concatenation $H * G$ (this is H followed by G) is a homotopy from h to f_n . Thus $\text{lift}_x(f_n)$ can be obtained by lifting the path $t \mapsto G(x, t)$ starting at the point $\text{lift}_x(f)$. The lift of $t \mapsto G(x, t)$ is a geodesic of length at most ε , so $|\text{lift}_x(f_n) - \text{lift}_x(f)| < \varepsilon$. \square

We now get back to the space of conformal embeddings $\text{CEmb}(X, Y, h)$ and look at where x can go under such maps.

Definition 6.4. $\text{Blob}(x)$ is the image of $\text{CEmb}(X, Y, h)$ by lift_x .

We know that $\text{Blob}(x)$ is at most 1-dimensional when $\text{CEmb}(X, Y, h)$ contains a slit mapping. Moreover, the assumption that h is generic eliminates the possibility that $\text{Blob}(x)$ be a circle.

Lemma 6.5. *If $\text{CEmb}(X, Y, h)$ contains a slit mapping, then $\text{Blob}(x)$ is homeomorphic to a point or a compact interval.*

Proof. By Lemma 4.13, the evaluation map $\text{ev}_x : \text{CEmb}(X, Y, h) \rightarrow Y$ is a homeomorphism onto its image $V(x)$, and the latter is either a point, a compact interval, or a circle. Then $\pi_Y : \text{Blob}(x) \rightarrow V(x)$ is a homeomorphism with inverse $\text{lift}_x \circ \text{ev}_x^{-1}$. Suppose that $\text{CEmb}(X, Y, h)$ is homeomorphic to a circle. Then there are \mathbb{Z} many homotopically distinct homotopies from any $f \in \text{CEmb}(X, Y, h)$ to itself, which contradicts Lemma 6.1. \square

We assume henceforth that $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping. To recapitulate the hypotheses thus far: X and Y are finite hyperbolic, ∂X is non-empty, $h : X \rightarrow Y$ is generic, and $\text{CEmb}(X, Y, h)$ is non-empty and does not contain any slit mapping.

Since $\text{CEmb}(X, Y, h)$ is compact and connected, and since lift_x is continuous, $\text{Blob}(x)$ is compact and connected as well. Our goal is to show that $\text{Blob}(x)$ is homeomorphic to a closed disk. The strategy of the proof is to analyze the boundary of $\text{Blob}(x)$. We will show that every point in $\partial \text{Blob}(x)$ is attained by a special kind of map in $\text{CEmb}(X, Y, h)$ which we call a slit mapping rel x .

Definition 6.6. Let $f \in \text{Map}(X, Y, h)$. We say that f is a *Teichmüller embedding rel x* if the restriction $f : X \setminus x \rightarrow Y \setminus f(x)$ is a Teichmüller embedding. A *slit mapping rel x* is a Teichmüller embedding rel x which is conformal.

The distinction to make here is that the initial and terminal quadratic differentials of f are allowed to have simple poles at x and $f(x)$ respectively. Here is a useful construction for moving a point around on a surface.

Lemma 6.7. *Given a smooth path $\gamma : [0, 1] \rightarrow \mathbb{D}$, there exists a quasiconformal diffeomorphism $\Pi : Y \rightarrow Y$ and a homotopy H from id_Y to Π such that the lift of $t \mapsto H(\pi_Y(\gamma(0)), t)$ based at $\gamma(0)$ is equal to γ .*

Proof. One way to construct Π is to break up the path $\pi_Y \circ \gamma$ into finitely many simple subarcs γ_j and define Π as the composition of as many diffeomorphisms Π_j . Push $\partial/\partial x$ forward by γ_j , extend this to a compactly supported vector field V_j on Y , and define Π_j to be the time t_j -flow for V_j where t_j is the time required to travel between the two endpoints of γ_j . The resulting diffeomorphism Π is the identity outside a compact set and is therefore quasiconformal. By construction, the concatenation of flows is a homotopy H from the identity to Π such that $H(\pi_Y(\gamma(0)), t) = \pi_Y(\gamma(t))$. The lift of $\pi_Y \circ \gamma$ based at $\gamma(0)$ is of course equal to γ . \square

A diffeomorphism such as in the previous lemma is called a *point-pushing diffeomorphism*. If a point y is on the complement of $\text{Blob}(x)$, then by definition y is not attained by any conformal embedding. Using Ioffe's theorem, we see that y is attained by a Teichmüller embedding rel x .

Lemma 6.8. *Let y be in the complement of $\text{Blob}(x)$. There exists a unique Teichmüller embedding f rel x homotopic to h such that $\text{lift}_x(f) = y$.*

Proof. Let γ be any smooth path from the basepoint b over $h(x)$ to y and let Π be a point-pushing diffeomorphism along the path. By definition of $\text{Blob}(x)$, there does not exist any $g \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_x(g) = y$. In other words, the set

$$\text{CEmb}(X \setminus x, Y \setminus \pi_Y(y), \Pi \circ h)$$

is empty. On the other hand, for any $F \in \text{CEmb}(X, Y, h)$ the map $\Pi \circ F$ is quasiconformal. By Ioffe's theorem there exists a Teichmüller embedding $f : X \setminus x \rightarrow Y \setminus \pi_Y(y)$ homotopic to $\Pi \circ h$. Then f extends to a quasiconformal embedding $f : X \rightarrow Y$ homotopic to $\Pi \circ h$ and hence to h . Thus f is a Teichmüller embedding rel x . Moreover,

$$\text{lift}_x(f) = \text{lift}_x(\Pi \circ h) = \text{lift}_{h(x)}(\Pi) = y.$$

If g is another Teichmüller embedding rel x homotopic to h such that $\text{lift}_x(g) = y$, then g is homotopic to f rel g . By Theorem 3.7, f and g have the same dilatation and $g \circ f^{-1}$ is a slit mapping. Since $g \circ f^{-1}$ fixes the puncture $\pi_Y(y)$, it is the identity by Corollary 4.11. \square

We now show that boundary points are attained by slit mappings rel x .

Lemma 6.9. *Let y be on the boundary of $\text{Blob}(x)$. There is a unique $f \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_x(f) = y$. The map f is a slit mapping rel x . If φ and ψ are initial and terminal quadratic differentials for f rel x , then φ has a simple pole at x and ψ has a simple pole at $f(x)$.*

Proof. Since $\text{Blob}(x)$ is closed, y belongs to $\text{Blob}(x)$, and hence there exists an element f of $\text{CEmb}(X, Y, h)$ such that $\text{lift}_x(f) = y$. Let $y_n \in \mathbb{D} \setminus \text{Blob}(x)$ be such that $y_n \rightarrow y$ as $n \rightarrow \infty$. Let γ_n be the geodesic from y to y_n and let $\Pi_n : Y \rightarrow Y$ be the corresponding point-pushing diffeomorphism. We can construct Π_n in such a way that it converges to the identity in the C^1 norm, so that its dilatation converges to 1 as $n \rightarrow \infty$.

By the previous lemma, there exists a Teichmüller embedding f_n rel x homotopic to h such that $\text{lift}_x(f_n) = y_n$. Moreover, the sequence of marked surfaces $[Y \setminus \pi_Y(y_n), \Pi_n]$ converges to $[Y \setminus \pi_Y(y), \text{id}]$ in $\mathcal{T}^\#(Y \setminus \pi_Y(y))$. Therefore, there exists a Teichmüller embedding $g : X \setminus x \rightarrow Y \setminus \pi(y)$ homotopic to f by Lemma 5.5. Since f is conformal, its dilatation is less than or equal to the dilatation of g . Hence f is a slit mapping rel x by Ioffe's theorem. If there is some $g \in \text{CEmb}(X, Y, h)$ be such that $\text{lift}_x(g) = y$, then g is homotopic to f rel x . Thus g is a slit mapping rel x by Ioffe's theorem and $g \circ f^{-1}$ is the identity since it fixes a puncture.

Suppose that the terminal quadratic differential ψ does not have a pole at y . Then f is a slit mapping from X to Y (not rel x). But we assumed that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Therefore ψ has a simple pole at $f(x)$ and $\varphi = f^*\psi$ has a simple pole at x . \square

This motivates the following definition.

Definition 6.10. If y , f , φ , and ψ are as in the previous lemma, then we say that f , φ , and ψ realize y .

We point out that although f is unique, φ and ψ need not be unique. Nevertheless, the set of quadratic differentials realizing a given point y in the boundary of $\text{Blob}(x)$ is convex.

Lemma 6.11. *Suppose that $\psi_0, \psi_1 \in \mathcal{Q}^+(Y \setminus \pi_Y(y))$ realize $y \in \partial \text{Blob}(x)$. Then for every $\alpha, \beta > 0$ the quadratic differential $\alpha\psi_0 + \beta\psi_1$ belongs to $\mathcal{Q}^+(Y \setminus \pi_Y(y))$, realizes y , and has a simple pole at $\pi_Y(y)$.*

Proof. $\alpha\psi_0 + \beta\psi_1 \geq 0$ along any arc in $\widehat{Y} \setminus f(X)$, where f is the slit mapping rel x realizing y . If $\alpha\psi_0 + \beta\psi_1$ does not have a pole at $\pi_Y(y)$, then f is a slit mapping from X to Y , contradiction. \square

We will see that any $\psi \in \mathcal{Q}^+(Y \setminus \pi_Y(y))$ which realizes y tells us something about the shape of $\text{Blob}(x)$ near y . In order to explain this, we first need to discuss measured foliations and extremal length.

7. EXTREMAL LENGTH OF PARTIAL MEASURED FOLIATIONS

There are several ways to define extremal length for measured foliations. We follow the approach developed in [16] and [15]. Throughout this section, we use the expression *almost-smooth* to mean continuous and continuously differentiable except perhaps at finitely many points.

Definition 7.1. A *partial measured foliation* on a Riemann surface X is a collection of open sets $U_j \subset X$ together with almost-smooth functions $v_j : U_j \rightarrow \mathbb{R}$ satisfying

$$v_j = \pm v_k + c_{jk}$$

on $U_j \cap U_k$, where c_{jk} is locally constant. Since $|dv_j| = |dv_k|$ on $U_j \cap U_k$, there is a well-defined 1-form $|dv|$ on $\bigcup_j U_j$, called the *transverse measure*. We extend $|dv|$ to be zero on $X \setminus \bigcup_j U_j$. The *leaves* of $|dv|$ are the maximal connected subsets of $\bigcup_j U_j$ on which each v_j is locally constant.

Remark. We could relax the regularity condition on the functions v_j and only assume that they belong to the Sobolev space $W^{1,2}$. For the sake of brevity and simplicity we will stick to the almost-smooth condition.

For example, if φ is a quadratic differential on X and $U_j \subset X$ is a simply connected domain on which φ does not have any singularities, then the function

$$v_j(z) = \text{Im} \int_{z_0}^z \sqrt{\varphi}$$

is well-defined on U_j up to an additive constant and sign. The leaves of the resulting partial measured foliation are the horizontal trajectories of φ . We write $\mathcal{F}[\varphi]$ for the resulting partial measured foliation $|dv| = |\text{Im} \sqrt{\varphi}|$.

Definition 7.2. The *Dirichlet energy* of a partial measured foliation $|dv|$ on X is

$$\text{Dir}(|dv|) := \int |\nabla v|^2 dA$$

where the gradient ∇v (only defined up to sign) is computed with respect to any smooth conformal metric on X with corresponding area form dA .

Alternatively, we can write

$$\text{Dir}(|dv|) = \int \left(\frac{\partial v}{\partial x} \right)^2 + \left(\frac{\partial v}{\partial y} \right)^2 dx dy$$

where $z = x + iy$ is any conformal coordinate.

For example, if φ is a holomorphic quadratic differential on X and $\mathcal{F}[\varphi] = |dv| = |\text{Im } \sqrt{\varphi}|$ is its horizontal foliation, then

$$\text{Dir}(\mathcal{F}[\varphi]) = \int_X |\varphi| = \|\varphi\|,$$

as is easily seen by computing the Dirichlet energy in natural coordinates where $\varphi = dz^2$ and $v(z) = \pm \text{Im } z + c$.

Definition 7.3. A *cross-cut* is a path $\alpha : (0, 1) \rightarrow X$ which extends to a continuous path $\widehat{\alpha} : [0, 1] \rightarrow \widehat{X}$. Two cross-cuts are *homotopic* if there is a homotopy through cross-cuts between them. The endpoints of the extended paths are allowed to move on their respective boundary components during the homotopy.

Definition 7.4. The height of a homotopy class c of closed curves or cross-cuts in X with respect to a partial measured foliation $|dv|$ is the quantity

$$\text{height}(c; |dv|) := \inf_{\alpha \in c} \int_{\alpha} |dv|,$$

where the infimum is taken over all piecewise smooth curves $\alpha \in c$.

Definition 7.5. Given two partial measured foliations $|dv|$ and $|dw|$ on X , we say that $|dv| \geq |dw|$ if

$$\text{height}(c; |dv|) \geq \text{height}(c; |dw|)$$

for every homotopy class c of closed curves or cross-cuts in X . The partial measured foliations $|dv|$ and $|dw|$ are *measure equivalent* if $|dv| \geq |dw|$ and $|dw| \geq |dv|$, i.e. if

$$\text{height}(c; |dv|) = \text{height}(c; |dw|)$$

for every homotopy class c of closed curves or cross-cuts in X .

Definition 7.6. The *extremal length* of a partial measured foliation $|dv|$ on X is defined as $\text{EL}(|dv|; X) := \inf \{ \text{Dir}(|dw|) : |dw| \geq |dv| \text{ on } X \}$.

Note that the extremal length of $|dv|$ only depends on the measure equivalence class of $|dv|$. Moreover, extremal length rescales quadratically in the sense that $\text{EL}(\lambda|dv|; X) = \lambda^2 \text{EL}(|dv|; X)$ for every $\lambda > 0$.

The fundamental result about extremal length is that the horizontal foliation of any integrable holomorphic quadratic differential minimizes Dirichlet energy in its equivalence class.

Theorem 7.7 (Marden–Strebel, Gardiner–Lakic). *Let $\varphi \in \mathcal{Q}(X)$ and let $\mathcal{F}[\varphi] = |\text{Im } \sqrt{\varphi}|$. Then $\text{EL}(\mathcal{F}[\varphi]; X) = \text{Dir}(\mathcal{F}[\varphi]) = \|\varphi\|$.*

This was first proved for measured foliations in [24], then generalized in [13]. See chapters 2 and 11 in [14] or chapter VII in [37] for detailed expositions. The version for partial measured foliations appears in [16] and [15], but the proof from [24] applies verbatim.

The next definition justifies our preference for partial measured foliations over measured foliations. The push-forward of a partial measured foliation by a smooth embedding is naturally a partial measured foliation, whereas this is not so clear for measured foliations.

Definition 7.8. If $|dv|$ is a partial measured foliation on X and $f : X \rightarrow Y$ is an almost-smooth embedding, then the *push-forward* $f_*|dv|$ is given by the collection of open sets $f(U_j)$ and the almost-smooth functions $v_j \circ f^{-1} : f(U_j) \rightarrow \mathbb{R}$.

We now show that extremal length can increase by a factor at most K under K -quasiconformal embeddings. In particular, extremal length does not increase under conformal embeddings.

Lemma 7.9. *Let $\varphi \in \mathcal{Q}(X)$ and let $f : X \rightarrow Y$ be an almost-smooth K -quasiconformal embedding. Then $\text{EL}(f_*\mathcal{F}[\varphi]; Y) \leq K \text{EL}(\mathcal{F}[\varphi]; X)$.*

Proof. Let $\zeta = \sigma + i\tau$ be a conformal coordinate on $f(X)$ and let $z = x + iy$ be a conformal coordinate at $f^{-1}(\zeta)$. Since f is K -quasiconformal, we have

$$dx dy = \det(df^{-1}) d\sigma d\tau \geq K^{-1} \|df^{-1}\|^2 d\sigma d\tau.$$

Let $v(z) = \text{Im} \int_{z_0}^z \sqrt{\varphi}$ so that $|dv| = |\text{Im} \sqrt{\varphi}| = \mathcal{F}[\varphi]$. We compute

$$\begin{aligned} \text{Dir}(f_*|dv|) &= \int_{f(X)} |\nabla(v \circ f^{-1})(\zeta)|^2 d\sigma d\tau \\ &= \int_{f(X)} |(df^{-1})(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &\leq \int_{f(X)} \|df^{-1}\|^2 |(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &\leq K \int_{f(X)} \det(df^{-1}) |(\nabla v)(f^{-1}(\zeta))|^2 d\sigma d\tau \\ &= K \int_X |(\nabla v)(z)|^2 dx dy = K \text{Dir}(|dv|). \end{aligned}$$

It follows that

$$\text{EL}(f_*|dv|; Y) \leq \text{Dir}(f_*|dv|) \leq K \text{Dir}(|dv|) = K \text{EL}(|dv|; X),$$

where the last equality holds by Theorem 7.7. \square

The inequality is sharp as the case of Teichmüller embeddings illustrates.

Lemma 7.10. *Let $f : X \rightarrow Y$ be a Teichmüller embedding of dilatation K with initial and terminal quadratic differentials φ and ψ . Then $f_*\mathcal{F}[\varphi]$ is*

measure equivalent to $\mathcal{F}[\psi]$ on Y and we have

$$\text{EL}(\mathcal{F}[\psi]; Y) = \text{EL}(f_*\mathcal{F}[\varphi]; Y) = K \text{EL}(\mathcal{F}[\varphi]; X).$$

Proof. We have $\mathcal{F}[\psi] = f_*\mathcal{F}[\varphi]$ on $f(X)$ since $f(x + iy) = Kx + iy$ in natural coordinates. Moreover, as $Y \setminus f(X)$ is a finite union horizontal arcs, the integral of $\mathcal{F}[\psi]$ is zero along any piece of curve contained in $Y \setminus f(X)$. For any piecewise smooth curve α in Y , we thus have

$$\begin{aligned} \int_{\alpha} \mathcal{F}[\psi] &= \int_{\alpha} \chi_{f(X)} \mathcal{F}[\psi] + \int_{\alpha} \chi_{Y \setminus f(X)} \mathcal{F}[\psi] \\ &= \int_{\alpha} f_*\mathcal{F}[\varphi] + \int_{\alpha} \chi_{Y \setminus f(X)} \mathcal{F}[\psi] \\ &= \int_{\alpha} f_*\mathcal{F}[\varphi]. \end{aligned}$$

Therefore $\mathcal{F}[\psi]$ is measure equivalent to $f_*\mathcal{F}[\varphi]$ which by definition implies the equality $\text{EL}(\mathcal{F}[\psi]; Y) = \text{EL}(f_*\mathcal{F}[\varphi]; Y)$. By Lemma 7.9 the inequality

$$\text{EL}(f_*\mathcal{F}[\varphi]; Y) \leq K \text{EL}(\mathcal{F}[\varphi]; X)$$

holds and by Theorem 7.7, we have

$$\text{EL}(\mathcal{F}[\psi]; Y) = \|\psi\| \quad \text{and} \quad \text{EL}(\mathcal{F}[\varphi]; X) = \|\varphi\|.$$

Lastly, we have $\|\psi\| = K\|\varphi\|$ because f stretches horizontally by a factor K and $f(X)$ has full measure in Y . Putting everything together, we get

$$\|\psi\| = \text{EL}(\mathcal{F}[\psi]; Y) = \text{EL}(f_*\mathcal{F}[\varphi]; Y) \leq K \text{EL}(\mathcal{F}[\varphi]; X) = K\|\varphi\| = \|\psi\|$$

and hence equality holds. \square

We will need a sufficient condition for when the push-forwards of a partial measured foliation by two homotopic embeddings are measure equivalent.

Definition 7.11. We say that an embedding $f : X \rightarrow Y$ is *tame* if it is almost-smooth and extends to a continuous map $\hat{X} \rightarrow \hat{Y}$ which is piecewise smooth along ∂X .

For example, every Teichmüller embedding is tame.

Lemma 7.12. Let $\varphi \in \mathcal{Q}^+(X)$ and let $f, g : X \rightarrow Y$ be homotopic tame embeddings. Then $f_*\mathcal{F}[\varphi]$ and $g_*\mathcal{F}[\varphi]$ are measure equivalent on Y .

Proof. By symmetry, it suffices to show that $f_*\mathcal{F}[\varphi] \geq g_*\mathcal{F}[\varphi]$. Let α be a piecewise smooth closed curve or cross-cut in Y . Given $\varepsilon > 0$, we have to find a closed curve or cross-cut β homotopic to α such that

$$\int_{\beta} g_*\mathcal{F}[\varphi] \leq \int_{\alpha} f_*\mathcal{F}[\varphi] + \varepsilon.$$

Let n be the number of components of $\alpha \cap f(X)$. Let $\delta = \varepsilon/4(n+1)$ and let X^δ be X minus a δ -neighborhood of its ideal boundary in the metric $|\varphi|^{1/2}$.

The map $g \circ f^{-1} : f(X^\delta) \rightarrow g(X^\delta)$ can be extended to a diffeomorphism $H : Y \rightarrow Y$ homotopic to the identity (see [25, Lemma 2]).

We will take β to be a slightly modified version of $H(\alpha)$. We keep the part $H(\alpha) \cap g(X^\delta)$, which has at most n components, unchanged. Then $H(\alpha) \setminus g(X^\delta)$ has at most $n+1$ components, the $+1$ occurring if α is a cross-cut. We homotope each component c of $H(\alpha) \setminus g(X^\delta)$ within $Y \setminus g(X^\delta)$ to an arc γ_j with at most three parts: the beginning γ_j^- and end γ_j^+ contained in $g(X \setminus X^\delta)$, and perhaps a middle portion γ_j^0 contained in $Y \setminus g(X)$. The reason we can do this is that $Y \setminus g(X)$ is a deformation retract of $Y \setminus g(X^\delta)$. Then we homotope each of γ_j^\pm within $g(X \setminus X^\delta)$ to an arc η_j^\pm such that

$$\int_{\eta_j^\pm} g_* \mathcal{F}[\varphi] = \int_{g^{-1}(\eta_j^\pm)} \mathcal{F}[\varphi] < \frac{\varepsilon}{2(n+1)}.$$

This is possible since the height of each component of $X \setminus X^\delta$ is at most $\delta = \varepsilon/4(n+1)$. Moreover, no matter how many times $g^{-1}(\gamma_j^\pm)$ winds around the annular component of $X \setminus X^\delta$ in which it is contained, we can push this winding part toward the boundary of X which is horizontal for φ . In doing so, the height of the winding part tends to zero and is thus eventually less than $\varepsilon/4(n+1)$, for a total height of at most $\varepsilon/2(n+1)$.

We then let β be the concatenation of $H(\alpha) \cap g(X^\delta)$ with the arcs $\eta_j^- \cup \gamma_j^0 \cup \eta_j^+$. By construction β is homotopic to α and we compute

$$\begin{aligned} \int_\beta g_* \mathcal{F}[\varphi] &= \int_{H(\alpha) \cap g(X^\delta)} g_* \mathcal{F}[\varphi] + \sum_j \int_{\eta_j^-} g_* \mathcal{F}[\varphi] + \sum_j \int_{\eta_j^+} g_* \mathcal{F}[\varphi] \\ &\leq \int_{\alpha \cap f(X^\delta)} f_* \mathcal{F}[\varphi] + \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \leq \int_\alpha f_* \mathcal{F}[\varphi] + \varepsilon \end{aligned}$$

which is what we wanted. \square

8. THE BLOB IS SEMI-SMOOTH

In this section we show that the boundary of $\text{Blob}(x)$ satisfies a regularity condition which we call semi-smoothness. We need the following formula due to Gardiner for the derivative of extremal length [13].

Theorem 8.1 (Gardiner's formula). *Let Z_0 be a finite Riemann surface and let $q \in \mathcal{Q}(Z_0)$. If μ is a smooth Beltrami form on Z_0 with $\|\mu\|_\infty < 1$ and $f_\mu : Z_0 \rightarrow Z_\mu$ is the solution to the Beltrami equation, then*

$$\text{EL}((f_\mu)_* \mathcal{F}[q]; Z_\mu) = \|q\| + 2\langle \mu, q \rangle + o(\|\mu\|_\infty).$$

Remark. In his proof, Gardiner shows that

$$\log \text{EL}((f_\mu)_* \mathcal{F}[q]; Z_\mu) \geq \log \|q\| + \frac{2}{\|q\|} \langle \mu, q \rangle - c_1 \|\mu\|_\infty^2.$$

for a universal constant c_1 . We will use this inequality instead.

We apply this formula in the special case where $Z_\mu = Y \setminus \gamma(t)$ for some analytic path $\gamma : I \rightarrow Y$. The pairing $\langle \mu, q \rangle$ is then approximately proportional to the real part of the residue of q in the direction of $\gamma'(0)$. See [27] for a similar but more general calculation.

Lemma 8.2. *Let γ be an analytic arc on Y and let $Z_t = Y \setminus \gamma(t)$. Then $Z_t = Z_{\mu(t)}$ for a smooth Beltrami form $\mu(t)$ such that $\|\mu(t)\|_\infty = O(t)$ and for every $q \in \mathcal{Q}(Z_0)$ we have*

$$\langle \mu(t), q \rangle = -\pi t \operatorname{Re}[\operatorname{Res}_{\gamma(0)}(q\gamma'(0))] + O(t^2).$$

Proof. Extend $\gamma'(t)$ to a holomorphic vector field \mathbf{v} on a disk D centered at $y = \gamma(0)$ and let ϕ be a smooth bump function which vanishes outside D and is equal to 1 on a disk D' around y compactly contained in D . Then $\phi\mathbf{v}$ is a smooth vector field defined on all of Y . Let Φ_t be the time- t flow for $\phi\mathbf{v}$ and let $\mu(t) = \bar{\partial}\Phi_t/\partial\Phi_t$. We have $\mu'(0) = \bar{\partial}(\phi\mathbf{v})$ and hence $\mu(t) = t\bar{\partial}(\phi\mathbf{v}) + O(t^2)$. We compute

$$\begin{aligned} \int_{Z_0} \mu'(0)q &= \int_D q\bar{\partial}(\phi\mathbf{v}) = \int_{D \setminus D'} q\bar{\partial}(\phi\mathbf{v}) = \int_{D \setminus D'} \bar{\partial}(\phi q\mathbf{v}) \\ &= -\frac{i}{2} \int_{\partial(D \setminus D')} \phi q\mathbf{v} = \frac{i}{2} \int_{\partial D'} q\mathbf{v} = -\pi \operatorname{Res}_y(q\mathbf{v}). \end{aligned}$$

The equality from the first line to the second is by Green's theorem and the change of sign in the next equality comes from reversing orientation on the circle $\partial D'$. To conclude the proof, recall that $\langle \nu, q \rangle = \operatorname{Re} \int_{Z_0} \nu q$ by definition. \square

Remark. The constants in the big O notation are not universal, but depend continuously on the size of the tangent vector $\gamma'(0)$ and the injectivity radius of Y at $\gamma(0)$.

We set up some notation before stating the next result. Given a point $\tilde{y} \in \mathbb{D}$ and a tangent vector $\tilde{\mathbf{v}} \in T_{\tilde{y}}\mathbb{D}$, let $y = \pi_Y(\tilde{y})$ and let $\mathbf{v} = d\pi_Y(\tilde{\mathbf{v}})$. For small $t \in \mathbb{R}$, we let $\tilde{y}(\tilde{\mathbf{v}}, t) = \tilde{y} + t\tilde{\mathbf{v}}$ and $y(\mathbf{v}, t) = \pi_Y(\tilde{y}(\tilde{\mathbf{v}}, t))$. Finally, we let $\Pi(\mathbf{v}, t) : Y \rightarrow Y$ be a diffeomorphism pushing \tilde{y} to $\tilde{y}(\tilde{\mathbf{v}}, t)$ in the universal cover. For any $\psi \in \mathcal{Q}(Y \setminus y)$, we compute the variation of the extremal length of $\mathcal{F}[\psi]$ along the path $t \mapsto Y \setminus y(\mathbf{v}, t)$.

Corollary 8.3. *Let $\tilde{y} \in \mathbb{D}$, let $\tilde{\mathbf{v}} \in T_{\tilde{y}}\mathbb{D}$, and let $\psi \in \mathcal{Q}(Y \setminus y)$ have area 1. Then*

$$\log \operatorname{EL}(\Pi(\mathbf{v}, t)_* \mathcal{F}[\psi]; Y \setminus y(\mathbf{v}, t)) \geq -2\pi t \operatorname{Re}[\operatorname{Res}_y(\psi\mathbf{v})] + O(t^2).$$

Proof. The inequality follows immediately from Lemma 8.2 and the remark following Theorem 8.1. \square

If $\psi \in \mathcal{Q}(Y \setminus y)$ has a simple pole at y , then there exists a tangent vector $\mathbf{v} \in T_y Y$ (unique up to rescaling) such that $\operatorname{Res}_y(\psi\mathbf{v}) < 0$. We say that \mathbf{v} is

vertical for ψ . For example, if $\psi = \frac{1}{z}dz^2$ and $\mathbf{v} = -\frac{\partial}{\partial z}$, then

$$\text{Res}_0(\psi\mathbf{v}) = -\frac{1}{2\pi i} \oint \frac{1}{z} dz = -1.$$

For every $e^{i\theta} \in \partial\mathbb{D}$ we have $\text{Res}_y(\psi e^{i\theta}\mathbf{v}) = e^{i\theta} \text{Res}_y(\psi\mathbf{v})$. If \mathbf{v} is vertical for ψ , then $\text{Re}[\text{Res}_y(\psi e^{i\theta}\mathbf{v})] = \cos(\theta) \text{Res}_y(\psi\mathbf{v})$ is negative precisely when $\theta \in (-\frac{\pi}{2}, \frac{\pi}{2})$. Thus we may improve the previous corollary as follows.

Lemma 8.4. *Let $\theta_0 \in (0, \frac{\pi}{2})$ and let K be a compact set of pairs (\tilde{y}, ψ) where $\tilde{y} \in \mathbb{D}$ and $\psi \in \mathcal{Q}(Y \setminus y)$ has area 1 and a simple pole at y . There exists $\delta > 0$ such that for every $(\tilde{y}, \psi) \in K$ we have*

$$\log \text{EL}(\Pi(e^{i\theta}\mathbf{v}, t)_*\mathcal{F}[\psi]; Y \setminus y(e^{i\theta}\mathbf{v}, t)) > 0$$

for every $t \in (0, \delta)$ and every $\theta \in [-\theta_0, \theta_0]$ where $\tilde{\mathbf{v}}$ is the vertical vector for $\pi_Y^*\psi$ at \tilde{y} rescaled to have norm 1 with respect to the hyperbolic metric.

Proof. By Corollary 8.3, we have

$$\log \text{EL}(\Pi(e^{i\theta}\mathbf{v}, t)_*\mathcal{F}[\psi]; Y \setminus y(e^{i\theta}\mathbf{v}, t)) \geq -2\pi t \cos(\theta) \text{Res}_y(\psi\mathbf{v}) - c(\tilde{y})t^2$$

where $c(\tilde{y})$ is a positive constant depending continuously on \tilde{y} . The right hand side is positive provided that

$$0 < t < -2\pi \cos(\theta) \text{Res}(\psi\mathbf{v})/c(\tilde{y}).$$

This upper bound is bounded below by some $\delta > 0$ since it is positive and depends continuously on θ , \tilde{y} , and ψ , which vary inside a compact set. \square

To see how the above inequality might be useful in studying $\text{Blob}(x)$, recall that conformal embeddings do not increase extremal length.

Lemma 8.5. *Let $f : X \rightarrow Y$ be a tame embedding homotopic to h , let $y = f(x)$, and let $\varphi \in \mathcal{Q}^+(X \setminus x)$. Suppose that $\Pi : Y \setminus y \rightarrow Y \setminus y'$ is a diffeomorphism such that*

$$\text{EL}((\Pi \circ f)_*\mathcal{F}[\varphi]; Y \setminus y') > \text{EL}(\mathcal{F}[\varphi]; X \setminus x).$$

Then there is no conformal embedding homotopic to $\Pi \circ f : X \setminus x \rightarrow Y \setminus y'$. In particular $\text{lift}_x(\Pi \circ f) \notin \text{Blob}(x)$.

Proof. Suppose there is a conformal embedding g homotopic to $\Pi \circ f \text{ rel } x$. By Corollary 5.9 we may assume that g is tame. Then $g_*\mathcal{F}[\varphi]$ is measure equivalent to $(\Pi \circ f)_*\mathcal{F}[\varphi]$ by Lemma 7.12. Since g is conformal we have

$$\text{EL}((\Pi \circ f)_*\mathcal{F}[\varphi]; Y \setminus y') = \text{EL}(g_*\mathcal{F}[\varphi]; Y \setminus y') \leq \text{EL}(\mathcal{F}[\varphi]; X \setminus x)$$

by Lemma 7.9, which is a contradiction. Hence $\text{lift}_x(\Pi \circ f)$ does not belong to $\text{Blob}(x)$, as there is no conformal embedding homotopic to $\Pi \circ f \text{ rel } x$. \square

Using the above results, we will show that if $\tilde{\mathbf{v}}$ is vertical for $\pi_Y^*\psi$ at $\tilde{y} \in \partial \text{Blob}(x)$ where ψ realizes \tilde{y} , then $\tilde{\mathbf{v}}$ is normal to $\text{Blob}(x)$ in the sense that it points orthogonally away from it.

Definition 8.6. Given $\mathbf{v} \in T_z\mathbb{C}$, $\theta \in (0, \pi)$, and $\delta > 0$, we denote by $\nabla(\mathbf{v}, \theta, \delta)$ the open angular sector based at z with radius δ and angle θ on either side of \mathbf{v} . In symbols,

$$\nabla(\mathbf{v}, \theta, \delta) = \left\{ z + te^{i\phi} \frac{\mathbf{v}}{\|\mathbf{v}\|} : \phi \in (-\theta, \theta) \text{ and } t \in (0, \delta) \right\}.$$

The closure of $\nabla(\mathbf{v}, \theta, \delta)$ is denoted $\nabla[\mathbf{v}, \theta, \delta]$.

Definition 8.7. Let $B \subset \mathbb{C}$ be closed. A vector $\mathbf{v} \in T_z\mathbb{C}$ with $z \in \partial B$ is *normal to B* if $\mathbf{v} \neq 0$ and if there are angular sectors arbitrarily close to half-disks pointing in the direction of \mathbf{v} which are disjoint from B . More precisely, \mathbf{v} is normal to B if $\mathbf{v} \neq 0$ and if for every $\theta \in (0, \frac{\pi}{2})$, there exists a $\delta > 0$ such that $\nabla(\mathbf{v}, \theta, \delta) \cap B = \emptyset$.

Theorem 8.8. Let $\tilde{y} \in \partial \text{Blob}(x)$. Suppose that f realizes \tilde{y} with respect to $\psi \in \mathcal{Q}^+(Y \setminus y)$ and that $\tilde{\mathbf{v}} \in T_{\tilde{y}}\mathbb{D}$ is vertical for $\pi_Y^*\psi$. Then $\tilde{\mathbf{v}}$ is normal to $\text{Blob}(x)$.

Proof. Fix $\theta_0 \in (0, \frac{\pi}{2})$. Then there exists a $\delta > 0$ such that for every $\theta \in [-\theta_0, \theta_0]$ and every $t \in (0, \delta)$ we have

$$\text{EL}(\Pi(e^{i\theta}\mathbf{v}, t)_*\mathcal{F}[\psi]; Y \setminus y(e^{i\theta}\mathbf{v}, t)) > \text{EL}(\mathcal{F}[\psi]; Y \setminus y) = \|\psi\|$$

by Lemma 8.4. Since f is a slit mapping, $\mathcal{F}[\psi]$ is measure equivalent to $f_*\mathcal{F}[\varphi]$ where $\varphi = f^*\psi$. We also have $\text{EL}(\mathcal{F}[\psi]; Y \setminus y) = \text{EL}(\mathcal{F}[\varphi]; X \setminus x)$ so that

$$\text{EL}((\Pi(e^{i\theta}\mathbf{v}, t) \circ f)_*\mathcal{F}[\varphi]; Y \setminus y(e^{i\theta}\mathbf{v}, t)) > \text{EL}(\mathcal{F}[\varphi]; X \setminus x)$$

for every $\theta \in [-\theta_0, \theta_0]$ and every $t \in (0, \delta)$. By Lemma 8.5,

$$\text{lift}_x(\Pi(e^{i\theta}\mathbf{v}, t) \circ f) = \tilde{y}(e^{i\theta}\tilde{\mathbf{v}}, t) = \tilde{y} + te^{i\theta}\tilde{\mathbf{v}}$$

does not belong to $\text{Blob}(x)$. In other words, $\nabla(\tilde{\mathbf{v}}, \theta_0, \varepsilon)$ is disjoint from $\text{Blob}(x)$, where ε is equal to δ times the Euclidean norm of $\tilde{\mathbf{v}}$. Since the angle $\theta_0 \in (0, \frac{\pi}{2})$ was arbitrary, $\tilde{\mathbf{v}}$ is normal to $\text{Blob}(x)$. \square

As a corollary, we obtain the converse of Lemma 6.9.

Corollary 8.9. If $f : X \rightarrow Y$ is a slit mapping rel x homotopic to h , then $\text{lift}_x(f) \in \partial \text{Blob}(x)$.

Proof. Since f is conformal, $\text{lift}_x(f) \in \text{Blob}(x)$. Let ψ be a terminal quadratic differential for f rel x . If ψ does not have a simple pole at $y = f(x)$ then f is a slit mapping, contradicting the assumption that $\text{CEmb}(X, Y, h)$ does not contain any. Thus ψ has a simple pole at y so that there exists a tangent vector $\mathbf{v} \in T_y Y$ such that $\text{Res}(\psi\mathbf{v}) < 0$. Let $\tilde{\mathbf{v}}$ be the lift of \mathbf{v} based at $\text{lift}_x(f)$. We can apply the same reasoning as in the proof of Theorem 8.8 to conclude that $\text{lift}_x(f) + t\tilde{\mathbf{v}} \notin \text{Blob}(x)$ for all $t > 0$ small enough, and hence that $\text{lift}_x(f) \in \partial \text{Blob}(x)$. \square

We will show the converse of Theorem 8.8, namely that every vector normal to $\text{Blob}(x)$ is vertical with respect to some realizing quadratic differential. We need a few lemmas first.

In Section 2, we mentioned that for every $Z \in \mathcal{T}^\#(S)$, the vector space $\mathcal{Q}(Z)$ is isomorphic to the tangent space to $\mathcal{T}^\#(S)$ at Z . It is also a fact that the vector spaces $\mathcal{Q}(Z)$ glue together to form a vector bundle \mathcal{Q}_S isomorphic to the tangent bundle of $\mathcal{T}^\#(S)$ [19].

Lemma 8.10. *Let $K \subset \mathbb{D}$ be compact. The set of meromorphic quadratic differentials ψ on Y such that $\psi \in \mathcal{Q}^+(Y \setminus y)$ for some $\tilde{y} \in K$ and such that $\|\psi\| = 1$ is compact.*

Proof. Given a basepoint $\tilde{y}_0 \in K$, we obtain a continuous map $\Theta : K \rightarrow \mathcal{T}^\#(Y \setminus y_0)$ by sending $\tilde{y} \in K$ to the corresponding point-pushing diffeomorphism $\Pi : Y \setminus y_0 \rightarrow Y \setminus y$. For any $y \in Y$, the real vector space $\mathcal{Q}(Y \setminus y)$ has finite dimension and hence its unit sphere is compact. The restriction of the unit sphere bundle $\mathbb{P}\mathcal{Q}_{Y \setminus y_0}$ over the compact set $\Theta(K)$ is thus compact. Since $\mathcal{Q}^+(Y \setminus y)$ is closed in $\mathcal{Q}(Y \setminus y)$, the result follows. \square

One can also deduce this from [26, Theorem A.3.1]. Next, we need a lemma saying that a limit of terminal quadratic differentials is a terminal quadratic differential for the limiting Teichmüller embedding.

Lemma 8.11. *Let $f_n : X \rightarrow Y$ be a Teichmüller embedding rel x of dilatation K_n with terminal quadratic differential ψ_n normalized to have area 1. Suppose that K_n is bounded above, that $f_n \rightarrow f$ and that $\psi_n \rightarrow \psi$ locally uniformly on $Y \setminus f(x)$. Then f is a Teichmüller embedding rel x with terminal quadratic differential ψ .*

Proof. Let φ_n be the initial quadratic differential for f_n corresponding to ψ_n . We may assume that K_n converges to some K and that φ_n converges to some $\varphi \in \mathcal{Q}^+(X \setminus x)$ since its norm $1/K_n$ is bounded below. By Lemma 5.5, f is a Teichmüller embedding rel x . We have to show that ψ is the terminal quadratic differential of f corresponding to φ .

Suppose that $z_0 \in X \setminus x$ is not a zero of φ . Then there is a compact simply connected neighborhood U of z_0 on which φ does not vanish. If n is large enough, then φ_n does not have any zeros in U either. If $V = f(U)$, then ψ and ψ_n do not have zeros in V when n is large enough. We can choose square roots consistently so that $\sqrt{\varphi_n} \rightarrow \sqrt{\varphi}$ uniformly on U and $\sqrt{\psi_n} \rightarrow \sqrt{\psi}$ uniformly on V . Then for every $z \in U$ and every n we have

$$\int_{f_n(z_0)}^{f_n(z)} \sqrt{\psi_n} = K_n \operatorname{Re} \left(\int_{z_0}^z \sqrt{\varphi_n} \right) + i \operatorname{Im} \left(\int_{z_0}^z \sqrt{\varphi_n} \right)$$

since f_n is a Teichmüller embedding of dilatation K_n with respect to φ_n and ψ_n . Taking the limit as $n \rightarrow \infty$ we get

$$\int_{f(z_0)}^{f(z)} \sqrt{\psi} = K \operatorname{Re} \left(\int_{z_0}^z \sqrt{\varphi} \right) + i \operatorname{Im} \left(\int_{z_0}^z \sqrt{\varphi} \right)$$

which means that f is locally of Teichmüller form with respect to φ and ψ . Observe that $\|\psi\| = K\|\varphi\|$ since $\|\psi_n\| = K_n\|\varphi_n\|$ for every n . Therefore $\tilde{f}(X)$ has full measure in Y , which by Lemma 3.6 implies that $\hat{Y} \setminus f(X)$ is a finite union of points and horizontal arcs for ψ . \square

The following lemma shows that the set of vectors vertical for some realizing quadratic differential at a given point is convex.

Lemma 8.12. *Let $\tilde{y} \in \partial \text{Blob}(x)$ and suppose that $\psi_0, \psi_1 \in \mathcal{Q}^+(Y \setminus y)$ realize \tilde{y} . If \mathbf{v}_0 and \mathbf{v}_1 are vertical for ψ_0 and ψ_1 respectively at y , then there exist $\alpha, \beta > 0$ such that $\alpha\psi_0 + \beta\psi_1$ is vertical for $\alpha\psi_0 + \beta\psi_1$.*

Proof. Take $\alpha = -\frac{|\mathbf{v}_0|}{|\mathbf{v}_1|} \text{Res}_y(\psi_1 \mathbf{v}_1)$ and $\beta = -\frac{|\mathbf{v}_1|}{|\mathbf{v}_0|} \text{Res}_y(\psi_0 \mathbf{v}_0)$. A calculation shows that $\text{Res}_y((\alpha\psi_0 + \beta\psi_1)(\mathbf{v}_0 + \mathbf{v}_1)) \leq 0$. By Lemma 6.11, the quadratic differential $\alpha\psi_0 + \beta\psi_1$ has a simple pole at y . This implies that $\mathbf{v}_0 + \mathbf{v}_1 \neq 0$ and hence that $\text{Res}_y((\alpha\psi_0 + \beta\psi_1)(\mathbf{v}_0 + \mathbf{v}_1)) < 0$. \square

We are now able to show that normal vectors are vertical.

Theorem 8.13. *Let \mathbf{v} be normal to $\text{Blob}(x)$ at y . Then there exists a quadratic differential $\psi \in \mathcal{Q}^+(Y \setminus \pi_Y(y))$ realizing y such that \mathbf{v} is vertical for $\pi_Y^* \psi$.*

Proof. Let V_y denote the set of vectors which are vertical for some quadratic differential realizing y . By Lemmas 6.11 and 8.12, V_y is convex. Moreover $V_y \cup \{\mathbf{0}_y\}$ is closed. Suppose that \mathbf{v} is not in V_y . Then there is an open half-plane H through y containing V_y such that \mathbf{v} is not in the closure \overline{H} . Let y_n be a sequence converging to y along the ray r which is normal to \overline{H} at y . Since r makes an angle strictly less than $\frac{\pi}{2}$ with \mathbf{v} and since \mathbf{v} is normal to $\text{Blob}(x)$, we may assume that y_n is not in $\text{Blob}(x)$ for any n . Let f_n be the Teichmüller embedding rel x realizing y_n provided by Lemma 6.8 and let ψ_n be its terminal quadratic differential, normalized to have area 1.

We may assume that ψ_n converges to some $\psi \in \mathcal{Q}^+(Y \setminus \pi_Y(y_n))$ and that f_n converges to a slit mapping f rel x with terminal quadratic differential ψ . If ψ is holomorphic at $\pi_Y(y)$, then f is a slit mapping on X , contrary to the assumption that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Therefore ψ has a simple pole at $\pi_Y(y)$, which implies that ψ_n has a simple pole at $\pi_Y(y_n)$ for all but finitely many indices. Let \mathbf{w}_n be the vertical direction for $\pi_Y^* \psi_n$ at y_n . By rescaling, we may assume that \mathbf{w}_n converges to a non-zero tangent vector \mathbf{w} vertical for $\pi_Y^* \psi$ at y . We have $\mathbf{w} \in V_y \subset H$. We will see that this yields a contradiction.

Let $\phi \in (0, \frac{\pi}{2}]$ be the angle that \mathbf{w} makes with the line ∂H and let $\theta = \frac{\pi}{2} - \phi$. By Lemma 8.4 there exists a $\delta > 0$ such that for every n and every $w \in \nabla(\mathbf{w}_n, \theta, \delta)$ we have

$$\text{EL}(\Pi(y_n \rightarrow w)_* \mathcal{F}[\psi_n]; Y \setminus \pi_Y(w)) > \text{EL}(\mathcal{F}[\psi_n]; Y \setminus \pi_Y(y_n))$$

On the other hand, since $\mathbf{w}_n \rightarrow \mathbf{w}$, the angle between \mathbf{w}_n and $y - y_n$ is eventually less than θ . Thus if n is large enough then $y \in \nabla(\mathbf{w}_n, \theta, \delta)$ and

hence

$$\text{EL}(\Pi(y_n \rightarrow y)_* \mathcal{F}[\psi_n]; Y \setminus \pi_Y(y)) > \text{EL}(\mathcal{F}[\psi_n]; Y \setminus \pi_Y(y_n)).$$




FIGURE 7. The point y is eventually contained in the sector of angle θ and radius δ about the vector \mathbf{w}_n .

Let φ_n be the initial quadratic differential of f_n corresponding to ψ_n and let K_n be the dilatation of f_n . By Lemma 7.10, $\mathcal{F}[\psi_n]$ is measure equivalent to $(f_n)_* \mathcal{F}[\varphi_n]$ on $Y \setminus \pi_Y(y_n)$. If we let $\Pi_n = \Pi(y_n \rightarrow y)$, then

$$\begin{aligned} \text{EL}((\Pi_n \circ f_n)_* \mathcal{F}[\varphi_n]; Y \setminus \pi_Y(y)) &= \text{EL}((\Pi_n)_* \mathcal{F}[\psi_n]; Y \setminus \pi_Y(y)) \\ &> \text{EL}(\mathcal{F}[\psi_n]; Y \setminus \pi_Y(y_n)) \\ &= K_n \text{EL}(\mathcal{F}[\varphi_n]; X \setminus x) \\ &> \text{EL}(\mathcal{F}[\varphi_n]; X \setminus x). \end{aligned}$$

By Lemma 8.5, it follows that $y = \text{lift}(\Pi_n \circ f_n)$ does not belong to $\text{Blob}(x)$, which is a contradiction. \square

We finally come to the main result of this section, which is that $\text{Blob}(x)$ is semi-smooth.

Definition 8.14. A closed subset $B \subset \mathbb{C}$ is *semi-smooth* if

- for every $z \in \partial B$, the set of normal vectors to B at z is non-empty and convex;
- any non-zero limit of normal vectors is normal.

Theorem 8.15. $\text{Blob}(x)$ is semi-smooth.

Proof. For every $y \in \partial \text{Blob}(x)$ the set of normal vectors to $\text{Blob}(x)$ at y coincides with the set V_y of vertical vectors for quadratic differentials realizing y . The set V_y is convex by Lemmas 6.11 and 8.12. Suppose that $y_n \in \partial \text{Blob}(x)$, that $y_n \rightarrow y$, that \mathbf{v}_n is vertical for ψ_n , and that $\mathbf{v}_n \rightarrow \mathbf{v} \neq 0$. Then we can rescale ψ_n and pass to a subsequence such that it converges to some ψ realizing y . We have $\text{Res}_y(\psi \mathbf{v}) \leq 0$ since $\text{Res}_{y_n}(\psi_n \mathbf{v}_n) < 0$ for every n . Moreover, ψ must have a simple pole at y for otherwise $\text{CEmb}(X, Y, h)$ would contain a slit mapping. This means that $\text{Res}_y(\psi \mathbf{v}) \neq 0$. Therefore \mathbf{v} is vertical for ψ hence normal to $\text{Blob}(x)$ at y . \square

9. THE BLOB IS A DISK

In this section, we show that the blob is homeomorphic to a closed disk. We first prove that every semi-smooth set is a manifold.

Theorem 9.1 (The aquatic theorem). *Every closed semi-smooth subset of \mathbb{C} is a 2-dimensional manifold with boundary.*

Proof. Let B be a closed semi-smooth set. Every interior point of B has a neighborhood homeomorphic to an open subset of \mathbb{C} , namely the interior of B . Thus we only have to show that every boundary point $z \in B$ has a neighborhood homeomorphic to a half disk. By applying an isometry of the plane, we may assume that $z = 0$ and that the vector i bisects the cone N_0 of vectors normal to B at 0. Let ϕ be half the angle of N_0 , let $\alpha = \phi + \frac{\pi}{2}$ and let $\beta = \pi - \alpha$.





FIGURE 8. The cone of normal vectors N_0 and the angles α and β .

Since B is semi-smooth, we have $\phi < \frac{\pi}{2}$ and hence $\beta > 0$. Moreover, for every $\theta \in (0, \alpha)$ there exists a $\delta > 0$ such that the open sector $\nabla(i, \theta, \delta)$ is disjoint from B . We now show the existence of closed sectors pointing downwards contained in B .

Claim 9.2. *For every $\theta \in (0, \beta)$, there exists a $\delta > 0$ such that the closed angular sector $\nabla[-i, \theta, \delta]$ is contained in B .*


Proof of Claim. Suppose not. Then there exists a $\theta \in (0, \beta)$ and a sequence $\delta_n \searrow 0$ for which the corresponding angular sector $S_n = \nabla[-i, \theta, \delta_n]$ intersects the complement of B for every n . Let D_n be a closed disk in S_n disjoint from B . Slide the center of D_n in a straight line towards 0 until the disk first hits B , and let D_n^* be the resulting disk. The intersection points of D_n^* with B all lie on the half of ∂D_n^* which is closest to 0. Let z_n be any point in this intersection. Then z_n is on the boundary of B and the unit vector \mathbf{v}_n pointing from z_n to the center of D_n^* is normal to B . Since S_n shrinks to 0, we have $z_n \rightarrow z$. Each vector \mathbf{v}_n makes an angle at most $\theta + \frac{\pi}{2}$ with the downward direction. Therefore, the vectors \mathbf{v}_n can only accumulate onto vectors forming an angle at least $\beta - \theta$ with the cone N_0 . This contradicts the hypothesis that every limit of normal vectors is normal. \square

FIGURE 9. Bubbles floating to the surface of B .

Let $\theta^+ \in (0, \alpha)$, let $S^+ = \nabla(i, \theta^+, \delta^+)$ be disjoint from B , let $\theta^- \in (0, \beta)$, and let $S^- = \nabla[-i, \theta^-, \delta^-]$ be contained in B . Let $I \subset S^+$ be a compact horizontal segment symmetric about the vertical line through 0 and lying entirely above S^- . We define a map $p : I \rightarrow \partial B$ as follows. For $z \in I$, let z fall straight down until it first hits B , and let $p(z)$ be this first hitting point. Note that $p(x + iy) = x + iq(x, y)$ for some function q so that p is injective.

Claim 9.3. *The map p is continuous on some subinterval $J \subset I$ centered at the midpoint of I .*

Proof of Claim. It is easy to see that p is continuous at the midpoint $p^{-1}(0)$. This is because p keeps the x -coordinate unchanged and moreover, $p(z)$ is below S^+ and above S^- . Thus the y -coordinate of $p(z)$ converges to 0 as $z \rightarrow p^{-1}(0)$.

FIGURE 10. The vertical projection p is akin to rain falling on the surface of B . The projection is continuous (in fact Lipschitz) by the squeeze theorem.

Let $0 < \varepsilon < \beta/2$. By semi-smoothness, there exists a $\delta > 0$ such that if $w \in \partial B$ is within distance δ of 0, then every vector in N_w is within angle $\phi + \varepsilon$ of the upward direction. For every $w \in \partial B$ with $|w| < \delta$, let \mathbf{v}_w be the bisector of N_w and let ϕ_w be half the angle of N_w . For every $\theta_w^+ \in (0, \phi_w + \frac{\pi}{2})$ there is an open sector $S_w^+ = \nabla(\mathbf{v}_w, \theta_w^+, \delta_w^+)$ disjoint from B by definition of

N_w . Since

$$\phi + \varepsilon = \alpha - \frac{\pi}{2} + \varepsilon < \alpha + \frac{\beta - \pi}{2} = \frac{\alpha}{2} < \frac{\pi}{2} \leq \phi_w + \frac{\pi}{2},$$

we may choose θ_w^+ so that S_w^+ contains the vertical direction in its span. By Claim 9.2, there is also a closed sector $S_w^- = \nabla[-\mathbf{v}_w, \theta_w^-, \delta_w^-]$ contained in B for every $\theta_w^- \in (0, \frac{\pi}{2} - \phi_w)$. The angle that $-\mathbf{v}_w$ makes with the downward direction is equal to the angle that \mathbf{v}_w makes with the vertical direction, which is at most $\phi + \varepsilon - \phi_w$ hence strictly less than $\frac{\pi}{2} - \phi_w$. Thus we may choose θ_w^- so that S_w^- contains the downward direction in its interior.

By continuity of p at $p^{-1}(0)$, there is a closed interval $J \subset I$ centered at $p^{-1}(0)$ such that $p(J)$ is contained in the ball of radius δ about 0. Let $z \in J$, let $w = p(z)$, and let S_w^+ and S_w^- be angular sectors as described in the previous paragraph. Also let $K \subset S_w^+$ be a compact horizontal segment crossing the vertical line through w and lying entirely above S_w^- . By construction, the vertical segment from z to w intersects B only at w . Since B is closed, we may assume that the rectangle with bottom edge K and upper edge $L \subset I$ is disjoint from B , by making K shorter if necessary. For every $\zeta \in L$, the image $p(\zeta)$ is thus squeezed between S_w^+ and S_w^- , so that p is continuous at z . \square

Thus $p(J)$ is the graph of a continuous function. Let $0 < \delta < |J|/2$. For every $z \in J$ with $|x| < \delta$, draw the open vertical segment of length 2δ centered at $p(z)$, and let U_δ be the union of those segments. The continuity of p implies that U_δ is open.

Claim 9.4. *If δ is small enough, then the component of $U_\delta \setminus p(J)$ above $p(J)$ is disjoint from B and the component below $p(J)$ is contained in B .*

Proof of Claim. If δ is small enough, then the component C^+ of $U_\delta \setminus p(J)$ above $p(J)$ lies below J itself. By definition of p , for every $z \in J$ the open vertical segment between z and $p(z)$ is disjoint from B , so that C^+ is disjoint from B .

For the component lying below $p(J)$, we use the same idea as in the proof of Claim 9.2. Suppose that the result does not hold. Then there is a sequence $\delta_n \searrow 0$ such that for every n , there is a closed disk D_n contained in the component of $U_{\delta_n} \setminus p(J)$ below $p(J)$. Slide the center of D_n upwards until the disk first hits B , and let D_n^* this hitting disk. Every intersection point of D_n^* with B is on the upper half of ∂D_n^* . Let z_n be any point in that intersection. Then z_n is on the boundary of B and the unit vector \mathbf{v}_n pointing from z_n towards the center of D_n^* is normal to B . As $n \rightarrow \infty$, we have $z_n \rightarrow 0$. Moreover, the vectors \mathbf{v}_n only accumulate onto vectors forming an angle at least $\frac{\pi}{2}$ with the upwards direction at 0, hence outside N_0 . This contradicts the semi-smoothness of B at 0. \square

By the claim, if δ is small enough then $U_\delta \cap B$ is equal to the union of $U_\delta \cap p(J)$ with the component of $U_\delta \setminus p(J)$ below $p(J)$. This neighborhood of

z in B is clearly homeomorphic to the rectangle $(-\delta, \delta) \times (0, \delta]$, which in turn is homeomorphic to a half disk. Thus B is a 2-manifold with boundary. \square

Since $\text{Blob}(x)$ is semi-smooth, it is a 2-manifold with boundary.

Corollary 9.5. *Suppose that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. Then $\text{Blob}(x)$ is a compact connected 2-manifold with boundary.*

In particular, the blob is the closure of its interior. We use this to show that the blob depends continuously on parameters. We first need to define a topology on closed subsets of a space.

Let S be a topological space and let $\text{CL}(S)$ be the hyperspace of closed subsets of S . The *Fell topology* on $\text{CL}(S)$ is the topology generated by neighborhoods of the form $N(K, \mathcal{U})$, where $K \subset S$ is compact, \mathcal{U} is a finite collection of open subsets of S , and $N(K, \mathcal{U})$ is the set of $A \in \text{CL}(S)$ such that $A \cap K = \emptyset$ and $A \cap U \neq \emptyset$ for every $U \in \mathcal{U}$.

Theorem 9.6 (Fell). *For any topological space S , the hyperspace $\text{CL}(S)$ is compact. If S is locally compact, then $\text{CL}(S)$ is Hausdorff.*

See [12] for a proof. If S is first-countable and Hausdorff, then a sequence $\{A_n\} \subset \text{CL}(S)$ converges to $A \in \text{CL}(S)$ if and only if

- for every $a \in A$, there exist $a_n \in A_n$ such that $a_n \rightarrow a$;
- for every sequence $\{a_n\}$ with $a_n \in A_n$, if $\{a_n\}$ accumulates onto $a \in S$, then $a \in A$.

We use the Fell topology on closed subsets of \mathbb{D} . To prove convergence, we mostly rely on Fell's compactness theorem and the above criterion for sequences.

It is fairly clear that the blob depends upper semi-continuously on parameters. The same holds for its boundary.

Lemma 9.7. *$\text{Blob}(x, X)$ and $\partial \text{Blob}(x, X)$ depend upper semi-continuously on the pair (X, x) . More precisely, suppose that $X_n \setminus x_n \rightarrow X \setminus x$ in $\mathcal{T}^\#(X \setminus x)$, that $\text{Blob}(x_n, X_n) \rightarrow A$ in $\text{CL}(\mathbb{D})$, and that $\partial \text{Blob}(x_n, X_n) \rightarrow B$ in $\text{CL}(\mathbb{D})$. Then $A \subset \text{Blob}(x, X)$ and $B \subset \partial \text{Blob}(x, X)$.*

Proof. Let $y \in A$. By hypothesis there exist $y_n \in \text{Blob}(x_n, X_n)$ such that $y_n \rightarrow y$. Let $f_n \in \text{CEmb}(X_n, Y, h)$ be such that $\text{lift}_{x_n}(f_n) = y_n$. Let $\sigma_n : X \setminus x \rightarrow X_n \setminus x_n$ be the Teichmüller map in the right homotopy class and let K_n be its dilatation (which converges to 1 by hypothesis). For any $K > 1$, if n is large enough then $f_n \circ \sigma_n : X \rightarrow Y$ is a K -quasiconformal embedding homotopic to h . The space of all K -quasiconformal embeddings from X to Y homotopic to h is compact by Lemma 3.2. After passing to a subsequence, we may thus assume that $f_n \circ \sigma_n$ converges to some K -quasiconformal embedding $f : X \setminus x \rightarrow Y$ in the same homotopy class. Since f is K -quasiconformal for every $K > 1$, it is conformal. Then

$$\text{lift}_x(f) = \lim_{n \rightarrow \infty} \text{lift}_x(f_n \circ \sigma_n) = \lim_{n \rightarrow \infty} \text{lift}_{x_n}(f_n) = \lim_{n \rightarrow \infty} y_n = y,$$

so that $y \in \text{Blob}(x, X)$.

Now let $y \in B$ and let $y_n \in \partial \text{Blob}(x_n, X_n)$ be such that $y_n \rightarrow y$. By Lemma 6.9, there exists a slit mapping $f_n \text{ rel } x_n$ from X_n to Y homotopic to h such that $\text{lift}_{x_n}(f_n) = y_n$. By Lemma 5.5, we can pass to a subsequence such that f_n converges to some slit mapping $f \text{ rel } x$ from X to Y . Then $\text{lift}_x(f) = y$ so that $y \in \text{Blob}(x, X)$. Moreover, $y \in \partial \text{Blob}(x, X)$ by Corollary 8.9. \square

We do not know if the blob moves continuously in general, but it does when there are no slit mappings at the limiting parameters.

Lemma 9.8. *Suppose that $\text{CEmb}(X, Y, h)$ does not contain any slit mapping. If $X_n \setminus x_n \rightarrow X \setminus x$ in $\mathcal{T}^\#(X \setminus x)$, then $\text{Blob}(x_n, X_n) \rightarrow \text{Blob}(x, X)$ and $\partial \text{Blob}(x_n, X_n) \rightarrow \partial \text{Blob}(x, X)$ in $\text{CL}(\mathbb{D})$.*

Proof. By compactness of $\text{CL}(\mathbb{D})$, it suffices to prove that if $\text{Blob}(x_n, X_n)$ converges to A and $\partial \text{Blob}(x_n, X_n)$ converges to B , then $A = \text{Blob}(x, X)$ and $B = \partial \text{Blob}(x, X)$.

Let us prove convergence of the blobs first. By Lemma 9.7, we have $A \subset \text{Blob}(x, X)$. We claim that the interior of $\text{Blob}(x, X)$ is contained in A . Let y be in the interior of $\text{Blob}(x, X)$ and suppose that there is an infinite set $J \subset \mathbb{N}$ such that y is not contained in $\text{Blob}(x_n, X_n)$ for every $n \in J$. Then for every $n \in J$, there exists a Teichmüller embedding $f_n \text{ rel } x_n$ with $\text{lift}_{x_n}(f_n) = y$. After passing to a subsequence in J , we get that $f_n \rightarrow f$ for some Teichmüller embedding $f \text{ rel } x$ by Lemma 5.5. We have $\text{lift}_x(f) = y$ by continuity of lift_x . By Corollary 8.9, y is in the complement of the interior of $\text{Blob}(x, X)$. This is a contradiction, which means that y is contained in $\text{Blob}(x_n, X_n)$ for all but finitely many indices, and hence $y \in A$. Since A is closed and $\text{Blob}(x, X)$ is the closure of its interior, we have $\text{Blob}(x, X) \subset A$ and hence $A = \text{Blob}(x, X)$.

By Lemma 9.7, we have $B \subset \partial \text{Blob}(x, X)$. Let $y \in \partial \text{Blob}(x, X)$. Let U be any connected neighborhood of y . We claim that if n is large enough, then U intersects both the complement of $\text{Blob}(x_n, X_n)$ and the interior of $\text{Blob}(x_n, X_n)$. Suppose on the contrary that U is contained in $\text{Blob}(x_n, X_n)$ for every n in an infinite set $J \subset \mathbb{N}$. Then $U \subset A = \text{Blob}(x, X)$, which is nonsense since y is on the boundary of $\text{Blob}(x, X)$. Similarly, suppose that U is contained in the complement of $\text{Blob}(x_n, X_n)$ for every n in an infinite set $J \subset \mathbb{N}$. Then for every $z \in U$ and every $n \in J$ there is a Teichmüller embedding $f_n : X_n \rightarrow Y \text{ rel } x_n$ homotopic to h such that $\text{lift}_{x_n}(f_n) = z$. By Lemma 5.5, f_n converges to a Teichmüller embedding $f \text{ rel } x$ after passing to a subsequence. Then $\text{lift}_x(f) = z$ so that $z \in \partial \text{Blob}(x)$ by Corollary 8.9. This is a contradiction, which proves the claim. Let n be large enough so that U intersects both the interior and the complement of $\text{Blob}(x_n, X_n)$. Since U is connected, it also intersects $\partial \text{Blob}(x_n, X_n)$. Since U can be chosen arbitrarily small, this shows that $y \in B$. \square

Similarly, nested families of blobs move continuously. In what follows, the surface X_r is obtained from X by gluing a cylinder of modulus r to each ideal boundary component as in Section 5.

Lemma 9.9. *The maps $r \mapsto \text{Blob}(x, X_r)$ and $r \mapsto \partial \text{Blob}(x, X_r)$ are continuous on $[0, R]$.*

Proof. If $r \in [0, R)$ and $\rho \rightarrow r$, then $\text{Blob}(x, X_\rho) \rightarrow \text{Blob}(x, X_r)$ and $\partial \text{Blob}(x, X_\rho) \rightarrow \partial \text{Blob}(x, X_r)$ by Lemma 9.8, since $\text{CEmb}(X_r, Y, h)$ does not contain any slit mapping. It remains to prove continuity at $r = R$. By compactness of $\text{CL}(\mathbb{D})$ and Lemma 9.7, it suffices to show that if $r_n \nearrow R$, if $\text{Blob}(x, X_{r_n}) \rightarrow A$, and if $\partial \text{Blob}(x, X_{r_n}) \rightarrow B$, then $A \supset \text{Blob}(x, X_R)$ and $B \supset \partial \text{Blob}(x, X_R)$. Let $y \in \text{Blob}(x, X_R) = \partial \text{Blob}(x, X_R)$. Then $y \in \text{Blob}(x, X_{r_n})$ for every n . Indeed, $r_n \leq R$ means that $X_{r_n} \subset X_R$ and hence $\text{CEmb}(X_R, Y, h) \subset \text{CEmb}(X_{r_n}, Y, h)$. It follows that $y \in A$. Let U be a connected neighborhood of y . Then U intersects $\text{Blob}(x, X_{r_n})$ since $y \in \text{Blob}(x, X_{r_n})$. Thus U intersects the interior of $\text{Blob}(x, X_{r_n})$ because $\text{Blob}(x, X_{r_n})$ is the closure of its interior. Suppose that U is contained in $\text{Blob}(x, X_{r_n})$ for every n in an infinite set $J \subset \mathbb{N}$. Then U is contained in A and hence in $\text{Blob}(x, X_R)$. This is absurd since $\text{Blob}(x, X_R)$ has empty interior. Thus U intersects $\partial \text{Blob}(x, X_{r_n})$ for all large enough n and hence $y \in B$. \square

We use continuity to show that the blob has no holes and is thus homeomorphic to a closed disk.

Theorem 9.10. *$\text{Blob}(x)$ is homeomorphic to a closed disk.*

Proof. We will show that the complement of $\text{Blob}(x)$ is connected, which is sufficient. Let $z_1, z_2 \in \mathbb{D} \setminus \text{Blob}(x)$. Note that z_1 and z_2 are contained in $\mathbb{D} \setminus \text{Blob}(x, X_\rho)$ for every $\rho \in [0, R]$ as the blobs are nested. Let r be the infimum of the set of $\rho \in [0, R]$ such that z_1 and z_2 are in the same component of $\mathbb{D} \setminus \text{Blob}(x, X_\rho)$. The set of such ρ is non-empty since $\text{Blob}(x, X_R)$ is a point or a compact interval, and hence has connected complement.

Suppose that z_1 and z_2 belong to different components of $\mathbb{D} \setminus \text{Blob}(x, X_r)$. Then $r < R$. In particular, $\text{Blob}(x, X_r)$ is a 2-manifold and each boundary component of $\text{Blob}(x, X_r)$ is a simple closed curve. Let C_1 be the component of $\partial \text{Blob}(x, X_r)$ surrounding z_1 , let C_2 be the one surrounding z_2 , and let γ be a simple closed curve in the interior of $\text{Blob}(x, X_r)$ separating C_1 from C_2 . For all ρ close enough to r we have that $\partial \text{Blob}(x, X_\rho)$ is disjoint from γ . On the other hand, there is a sequence $\rho_n \searrow r$ such that z_1 and z_2 belong to the same component of $\mathbb{D} \setminus \text{Blob}(x, X_{\rho_n})$. Let γ_n be a path in $\mathbb{D} \setminus \text{Blob}(x, X_{\rho_n})$ connecting z_1 and z_2 . For every n , γ_n intersects γ , say at w_n . Since γ is compact, we may pass to a subsequence so that $w_n \rightarrow w$ for some $w \in \gamma$. Now w is in the interior of $\text{Blob}(x, X_r)$. Let U be an open disk centered at w whose closure is contained in the interior of $\text{Blob}(x, X_r)$. Since $\text{Blob}(x, X_{\rho_n}) \rightarrow \text{Blob}(x, X_r)$, the open set U must intersect $\text{Blob}(x, X_{\rho_n})$ for all large enough n . Since $w_n \in \gamma \setminus \text{Blob}(x, X_{\rho_n})$ and since $\gamma \cup U$ is connected,

the intersection of $\gamma \cup U$ with $\partial \text{Blob}(x, X_{\rho_n})$ is non-empty. Let ζ_n be in the intersection. After passing to a subsequence, ζ_n converges to some point ζ in $\gamma \cup \overline{U}$. This is a contradiction since $\partial \text{Blob}(x, X_{\rho_n}) \rightarrow \partial(x, X_r)$ but $\gamma \cup \overline{U}$ is disjoint from $\partial(x, X_r)$. Therefore z_1 and z_2 belong to the same component of $\mathbb{D} \setminus \text{Blob}(x, X_r)$.

Suppose that $r > 0$. Let γ be a path joining z_1 to z_2 in $\mathbb{D} \setminus \text{Blob}(x, X_r)$. Since γ is compact and $\text{Blob}(x, X_\rho)$ depends continuously on ρ , the two are disjoint for all ρ sufficiently close to r . Then z_1 and z_2 belong to the same component of $\mathbb{D} \setminus \text{Blob}(x, X_\rho)$ for all $\rho < r$ sufficiently close to r , which contradicts the minimality of r . We conclude that $r = 0$ and that z_1 and z_2 belong to the same component of $\mathbb{D} \setminus \text{Blob}(x, X_0)$. Since z_1 and z_2 were arbitrary, the complement of $\text{Blob}(x)$ is connected and thus $\text{Blob}(x)$ is homeomorphic to a closed disk. \square

10. THE DEFORMATION RETRACTION

Fix once and for all a countable dense subset $\{x_1, x_2, \dots\} \subset X$. For each $n \in \mathbb{N}$ choose a lift b_n of $h(x_n)$ to \mathbb{D} and define the map $\text{lift}_n = \text{lift}_{x_n}$ as in Section 6. Also fix a conformal embedding $F \in \text{CEmb}(X, Y, h)$ which maximizes \mathbf{m} . In this section, we construct a (strong) deformation retraction of $\text{CEmb}(X, Y, h)$ into $\{F\}$.

Given any $f \in \text{CEmb}(X, Y, h)$, we define a sequence of paths $\gamma_n : [0, 1] \rightarrow \mathbb{D}$ inductively as follows. Let $G[1] : \mathbb{D} \rightarrow \text{Blob}(x_1)$ be the Riemann map with $G[1](0) = \text{lift}_1(F)$ and $G[1]'(0) > 0$, and let

$$\gamma_1(t) = \begin{cases} \text{lift}_1(f) & \text{if } t \in [0, 1/2] \\ G[1]((2 - 2t)G[1]^{-1}(\text{lift}_1(f))) & \text{if } t \in [1/2, 1]. \end{cases}$$

In words, γ_1 stays at $\text{lift}_1(f)$ for half the time and then moves at constant speed along the conformal ray towards the “center” $\text{lift}_1(F)$ of $\text{Blob}(x_1)$. In particular, $\gamma_1(t)$ belongs to $\text{Blob}(x_1)$ for every $t \in [0, 1]$ so that there exists some $g \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_1(g) = \gamma_1(t)$.

Let $n \geq 2$. Suppose that paths $\gamma_1, \dots, \gamma_{n-1}$ have been defined in such a way that

- the points $\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_{n-1}(t))$ are distinct for every $t \in [0, 1]$;
- γ_j is constant on the interval $[0, 2^{-j}]$ for every $j \in \{1, \dots, n-1\}$;
- $\gamma_j(0) = \text{lift}_j(f)$ and $\gamma_j(1) = \text{lift}_j(F)$ for every $j \in \{1, \dots, n-1\}$.

Then let

$$X[n] = X \setminus \{x_1, \dots, x_{n-1}\}, \quad Y[n, t] = Y \setminus \{\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_{n-1}(t))\},$$

and let $h[n, t] = \Pi[n, t] \circ h$ where $\Pi[n, t] : Y \rightarrow Y$ is a point-pushing diffeomorphism such that $\text{lift}_j(\Pi[n, t]) = \gamma_j(t)$ for every $j \in \{1, \dots, n-1\}$. Also define

$$E[n, t] = \text{CEmb}(X[n], Y[n, t], h[n, t])$$

and

$$\text{Blob}[n, t] = \text{Blob}(x_n, X[n], Y[n, t], h[n, t]).$$

We assume that $E[n, t]$ is non-empty as part of the induction hypothesis. Note that $\text{Blob}[n, t]$ is either a closed disk or a point. Indeed, if $E[n, t]$ contains a slit mapping ($\text{rel } \{x_1, \dots, x_{n-1}\}$) then $\text{Blob}[n, t]$ is homeomorphic to a point or an interval by Lemma 4.13. But since every map in $E[n, t]$ sends the puncture x_1 to the puncture $\gamma_1(t)$, there is at most one slit mapping in $E[n, t]$ by Corollary 4.11. Also, since we chose the paths $\gamma_1, \dots, \gamma_{n-1}$ to be constant on $[0, 2^{1-n}]$, the set $\text{Blob}[n, t]$ does not change for t in that interval. The next step is to choose a conformal center for $\text{Blob}[n, t]$.

Lemma 10.1. *For every $t \in [0, 1]$ there is a unique map $g[n, t]$ maximizing \mathbf{m} within $E[n, t]$. The map $t \mapsto g[n, t]$ is continuous, constant on $[0, 2^{1-n}]$, and satisfies $g[n, 1] = F$.*

Proof. The map \mathbf{m} is upper semi-continuous on the compact space $E[n, t]$. It thus attains its maximum at some $g[n, t]$ say with value R . By Theorem 5.6, the maximal extension of $g[n, t]$ is a slit mapping from $X_R \setminus \{x_1, \dots, x_{n-1}\}$ to $Y[n, t]$. By Corollary 4.11, the map $g[n, t]$ is unique since it sends a puncture to a puncture. Any limit of the maximal extension of $g[n, t]$ as $t \rightarrow s$ is a slit mapping by Lemma 5.5 and thus its restriction to $X[n]$ maximizes \mathbf{m} in $E[n, s]$ by Lemma 5.10. Thus $g[n, t] \rightarrow g[n, s]$ as $t \rightarrow s$. The paths γ_j for $j \in \{1, \dots, n-1\}$ are all constant on $[0, 2^{1-n}]$ so $g[n, t]$ does not change on that interval. Finally, since F maximizes \mathbf{m} on $\text{CEmb}(X, Y, h)$, it maximizes \mathbf{m} on the subset $E[n, 1]$ as well, and we have $g[n, 1] = F$. \square

Let $G[n] : \overline{\mathbb{D}} \rightarrow \text{Blob}[n, 0]$ be the Riemann map normalized so that $G[n](0) = \text{lift}_n(g[n, 0])$ and $G[n]'(0) > 0$. Then let

$$\gamma_n(t) = \begin{cases} \text{lift}_n(f) & \text{if } t \in [0, 2^{-n}) \\ G[n]((2 - 2^n t)G[n]^{-1}(\text{lift}_n(f))) & \text{if } t \in [2^{-n}, 2^{1-n}) \\ \text{lift}_n(g[n, t]) & \text{if } t \in [2^{1-n}, 1]. \end{cases}$$

This means that γ_n stays put at $\text{lift}_n(f)$ for some time, then travels along the conformal ray towards the center of $\text{Blob}[n, 0] = \text{Blob}[n, 2^{1-n}]$, and then follows the center for the rest of the time. It is possible that $\text{Blob}[n, 0]$ is a point if $\text{lift}_{n-1}(f)$ is in the boundary of $\text{Blob}[n-1, 0]$. In that case we let $G[n] : \overline{\mathbb{D}} \rightarrow \text{Blob}[n, 0]$ be the constant map. In other words we simply keep γ_n constant on $[0, 2^{1-n}]$. By construction we have $\pi_Y(\gamma_n(t)) \in Y[n, t]$ which means that the points $\pi_Y(\gamma_1(t)), \dots, \pi_Y(\gamma_n(t))$ are distinct for every $t \in [0, 1]$. Moreover the path γ_n is constant on the interval $[0, 2^{-n}]$. Finally, $E[n+1, t]$ is non-empty since $\gamma_n(t) \in \text{Blob}[n, t]$ for every $t \in [0, 1]$. This finishes the induction scheme.

We now show that the paths $\{\gamma_n\}$ automatically define a path from f to F inside the space $\text{CEmb}(X, Y, h)$.

Lemma 10.2. *For every $t \in [0, 1]$, there exists a unique $f_t \in \text{CEmb}(X, Y, h)$ such that $\text{lift}_n(f_t) = \gamma_n(t)$ for every $n \in \mathbb{N}$. The map $t \mapsto f_t$ is continuous and satisfies $f_0 = f$ and $f_1 = F$.*

Proof. Observe that $E[n, t]$ is a non-empty closed subset of $\text{CEmb}(X, Y, h)$ and is thus compact. Therefore, for each $t \in [0, 1]$, the nested intersection $\bigcap_{n=1}^{\infty} E[n, t]$ is non-empty. Any two functions in the intersection agree on the dense set $\{x_1, x_2, \dots\}$ and hence on all of X . Therefore, there is a unique function f_t in the intersection. Moreover, f_t varies continuously with t . Indeed, if g is any limit of any subsequence of f_t as $t \rightarrow s$, then for every $n \in \mathbb{N}$, we have

$$g(x_n) = \lim_{t \rightarrow s} f_t(x_n) = \lim_{t \rightarrow s} \pi_Y \circ \gamma_n(t) = \pi_Y \circ \gamma_n(s) = f_s(x_n)$$

so that $g = f_s$. It follows that $f_t \rightarrow f_s$ as $t \rightarrow s$. By construction we have $\text{lift}_n(f) = \gamma_n(0)$ and $\text{lift}_n(F) = \gamma_n(1)$ for every $n \in \mathbb{N}$. \square

We thus have a map $H : \text{CEmb}(X, Y, h) \times [0, 1] \rightarrow \text{CEmb}(X, Y, h)$ defined by $H(f, t) = f_t$. This map is such that

- $t \mapsto H(f, t)$ is continuous for every $f \in \text{CEmb}(X, Y, h)$;
- $H(f, 0) = f$ and $H(f, 1) = F$ for every $f \in \text{CEmb}(X, Y, h)$;
- $H(F, t) = F$ for every $t \in [0, 1]$.

The last point holds because if $f = F$, then every path γ_n is constant and hence $f_t = F$ for every t . It remains to prove that H is continuous in both variables.

Lemma 10.3. *Suppose that for every $n \in \mathbb{N}$, the map $f \mapsto \gamma_n$ is continuous. Then H is continuous.*

Proof. If for every $n \in \mathbb{N}$ the map $(f, t) \mapsto H(f, t)(x_n)$ is continuous, then H is continuous. This is because of the compactness of $\text{CEmb}(X, Y, h)$ and the fact that $\{x_1, x_2, \dots\}$ is dense in X (see the proof of Lemma 10.2). Since $H(f, t)(x_n) = f_t(x_n) = \pi_Y(\gamma_n(t))$, it thus suffices to show that the map $(f, t) \mapsto \gamma_n(t)$ is continuous. This condition is equivalent to the requirement that $f \mapsto \gamma_n$ be continuous, where the space of continuous maps $[0, 1] \rightarrow \mathbb{D}$ is equipped with the compact open topology (which is the same as the topology of uniform convergence). This is because the interval $[0, 1]$ is locally compact Hausdorff (see [29, p.287]). \square

Since the map $f \mapsto \text{lift}_1(f)$ is continuous and the Riemann map $G[1]$ is continuous, it is easy to see that $f \mapsto \gamma_1$ is continuous. We proceed by induction for the rest. Let $n \geq 2$ and suppose that the maps $f \mapsto \gamma_j$ are all continuous for $j = 1, \dots, n-1$. We will prove that the map $f \mapsto G[n]$ is continuous, which obviously implies that $f \mapsto \gamma_n$ is continuous. We use the following theorem of Radó, a proof of which is given in [30, p.26].

Theorem 10.4 (Radó). *Let (D_k, w_k) and (D, w) be topological closed disks in \mathbb{C} , each with a marked point in the interior. Suppose that $w_k \rightarrow w$ and that $D_k \rightarrow D$ in the Fell topology. Suppose also that there are parametrizations $c_k : S^1 \rightarrow \partial D_k$ and $c : S^1 \rightarrow \partial D$ such that $c_k \rightarrow c$ uniformly. Then the normalized Riemann map $(\overline{\mathbb{D}}, 0) \rightarrow (D_k, w_k)$ converges uniformly on $\overline{\mathbb{D}}$ to the normalized Riemann map $(\overline{\mathbb{D}}, 0) \rightarrow (D, w)$.*

By a slight generalization⁶ of Lemma 9.7, the maps $(f, t) \mapsto \text{Blob}[n, t]$ and $(f, t) \mapsto \partial \text{Blob}[n, t]$ are upper semi-continuous in t . Moreover, they are continuous at every (f, t) for which $E[n, t]$ does not contain any slit mapping by Lemma 9.8. But if $E[n, t]$ contains a slit mapping, then $\text{Blob}[n, t] = \partial \text{Blob}[n, t]$ is a single point and thus upper semi-continuity at (f, t) implies continuity. By Lemma 10.1, the conformal center $\text{lift}_n(g[n, t])$ of $\text{Blob}[n, t]$ also depends continuously on (f, t) . The only thing that remains to be checked is that the boundary $\partial \text{Blob}[n, t]$ can be parametrized as to converge uniformly.

Definition 10.5. Let $\{c_k\}_{k=1}^\infty$ be a sequence of simple closed curves in $\widehat{\mathbb{C}}$. We say that $\{c_k\}_{k=1}^\infty$ has a *collapsing finger* if after passing to a subsequence, there exist $x_k, y_k, z_k, w_k \in S^1$ in cyclic order and $x, y \in \widehat{\mathbb{C}}$ with $x \neq y$ such that $c_k(x_k) \rightarrow x$, $c_k(y_k) \rightarrow y$, $c_k(z_k) \rightarrow x$, and $c_k(w_k) \rightarrow y$.




FIGURE 11. A piece of a curve with a finger about to collapse.

We now show that collapsing fingers are the only obstructions to uniform convergence of simple closed curves.

Theorem 10.6. Let c_k and c be simple closed curves in $\widehat{\mathbb{C}}$ such that $c_k(S^1)$ converges to $c(S^1)$ in the Fell topology. If $\{c_k\}_{k=1}^\infty$ does not have any collapsing finger, then we can reparametrize c_k such that $c_k \rightarrow c$ uniformly.

Proof. By the Jordan–Schoenflies Theorem, c can be extended to a homeomorphism $\hat{c} : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$. Then $\hat{c}^{-1} \circ c_k(S^1) \rightarrow S^1$ in the Fell topology and the sequence $\{\hat{c}^{-1} \circ c_k\}_{k=1}^\infty$ does not have any collapsing finger. Moreover, if $\sigma_k : S^1 \rightarrow S^1$ is a homeomorphism such that $\hat{c}^{-1} \circ c_k \circ \sigma_k$ converges uniformly to the inclusion map $S^1 \hookrightarrow \mathbb{C}$, then $c_k \circ \sigma_k$ converges uniformly to c . We may thus assume that c is the inclusion map $S^1 \hookrightarrow \mathbb{C}$.

If k is large enough, then $c_k(S^1)$ is disjoint from 0 and ∞ . We claim that if k is large enough, then the winding number of c_k around the origin is ± 1 . Since c_k is simple, its winding number is either -1 , 0 , or 1 . Suppose the claim is false. Then after passing to a subsequence, the winding number of c_k is 0 for every k . Let $\arg(c_k) = c_k/|c_k|$ and let $\text{Arg}(c_k) : S^1 \rightarrow \mathbb{R}$ be a lift of $\arg(c_k)$ under the universal covering map $\mathbb{R} \rightarrow S^1$. This lift exists because the winding number is zero. Let $[a_k, b_k]$ be the image of $\text{Arg}(c_k)$. Since $c_k(S^1)$ converges to S^1 , it follows that the image $\arg(c_k)(S^1)$ converges

⁶The codomain $Y[n, t]$ is not fixed but depends continuously on (f, t) . The results of Section 9 generalize easily to this situation.

to S^1 as well, and hence $\liminf_{n \rightarrow \infty} b_k - a_k \geq 2\pi$. Thus if k is large enough, then $b_k - a_k > \pi$. Let x_k and z_k in S^1 be such that $\text{Arg}(c_k)(x_k) = a_k$ and $\text{Arg}(c_k)(z_k) = \min(b_k, a_k + 2\pi)$. Also let $y_k \in \overline{x_k z_k}$ and $w_k \in \overline{z_k x_k}$ be such that

$$\text{Arg}(c_k)(y_k) = \text{Arg}(c_k)(w_k) = a_k + \pi.$$

Since $c_k(S^1) \rightarrow S^1$, we may pass to a subsequence so that $c_k(x_k)$, $c_k(z_k)$, $c_k(y_k)$, and $c_k(w_k)$ converge to some x , y , z , and w in S^1 . Then $x = z$, $y = w$, and $x \neq y$, i.e. $\{c_k\}_{n=1}^\infty$ has a collapsing finger. This is a contradiction, which proves the claim.

If the winding number of c_k around the origin is -1 , then we reverse the parametrization so that it becomes $+1$. Let $\zeta_1^k < \zeta_2^k < \dots < \zeta_k^k$ be a partition of S^1 into k congruent arcs. Since c_k has winding number 1, we can find $\xi_1^k < \xi_2^k < \dots < \xi_k^k$ in S^1 such that $\arg(c_k)(\xi_j^k) = \zeta_j^k$ for every $j \in \{1, \dots, k\}$. Let $\sigma_k : S^1 \rightarrow S^1$ be any homeomorphism such that $\sigma_k(\zeta_j^k) = \xi_j^k$ for every $j \in \{1, \dots, k\}$. We claim that $c_k \circ \sigma_k$ converges uniformly to the inclusion map $c : S^1 \hookrightarrow \mathbb{C}$.

To simplify notation, we assume that c_k was parametrized correctly from the start, i.e. we assume that for every $k \gg 0$ and every $j \in \{1, \dots, k\}$, we have $\arg(c_k)(\zeta_j^k) = \zeta_j^k$. If c_k does not converge uniformly to c , then there exists an $\varepsilon > 0$ and an infinite set $J \subset \mathbb{N}$ such that for every $k \in J$, there exists a $y_k \in S^1$ such that $|c_k(y_k) - y_k| \geq \varepsilon$. Since S^1 is compact and $c_k(S^1) \rightarrow S^1$, we can pass to a subsequence such that $y_k \rightarrow x$ and $c_k(y_k) \rightarrow y$ for some x and y in S^1 . Note that $|y - x| \geq \varepsilon$ and in particular $y \neq x$. Let $j \in \{1, \dots, k\}$ be such that $\zeta_j^k \leq y_k < \zeta_{j+1}^k$, where we define $\zeta_{k+1}^k = \zeta_1^k$. Then let $x_k = \zeta_j^k$ and $z_k = \zeta_{j+1}^k$. Also let $w_k \in \{\zeta_1^k, \dots, \zeta_k^k\}$ be the closest point to y which comes after z_k but before x_k in the cyclic order on S^1 . We have $c_k(x_k) = x_k \rightarrow x$, $c_k(y_k) \rightarrow y$, $c_k(z_k) = z_k \rightarrow x$, and $c_k(w_k) = w_k \rightarrow y$. In other words, the sequence $\{c_k\}_{k=1}^\infty$ has a collapsing finger, which is a contradiction. Therefore, c_k converges uniformly to c . \square

To conclude the proof, we show that $\partial \text{Blob}[n, t]$ does not have any collapsing fingers. The reason for this is that the blobs $\text{Blob}[n, t]$ are uniformly semi-smooth, meaning that any non-zero limit of a sequence of vectors normal to some blob is normal to the limiting blob⁷. Now if there was a collapsing finger somewhere, then we would see two normal vectors pointing opposite to each other in the limit, which is forbidden by the definition of semi-smoothness.

Theorem 10.7. *Suppose that $(f_k, t_k) \rightarrow (f, t)$ in $\text{CEmb}(X, Y, h) \times [0, 1]$. Then $\partial \text{Blob}[n, t_k]$ converges to $\partial \text{Blob}[n, t]$ without collapsing fingers.*

Proof. Let $B_k = \text{Blob}[n, t_k]$, $B = \text{Blob}[n, t]$, $c_k = \partial B_k$, and $c = \partial B$. Suppose that after passing to a subsequence we can find $x_k, y_k, z_k, w_k \in c_k$ in cyclic

⁷The proof of this is a straightforward generalization of Theorem 8.15

order and $x, y \in c$ with $x \neq y$ such that $x_k, z_k \rightarrow x$ and $y_k, w_k \rightarrow y$. Rotate and translate the picture in such a way that $x = 0$ and that the upward direction i bisects the cone N_0 of vectors normal to B at 0.

By the proof of Theorem 9.1, there exists a rectangle Q centered at 0 with sides parallel to the coordinate axes such that $Q \cap c$ is the graph of a continuous function. Since $y \neq 0$, we can shrink Q so that it does not contain y . Let $\delta > 0$ be such that the vertical δ neighborhood U_δ of $Q \cap c$ is contained in Q . Then $Q \setminus U_\delta$ is compact and disjoint from c . Let k be large enough so that x_k and z_k are in Q , y_k and w_k are not in Q , and c_k is disjoint from $Q \setminus U_\delta$. Then of the three subarcs $\overline{x_k y_k}$, $\overline{y_k z_k}$, and $\overline{z_k w_k}$ of c_k , at least two must cross the same vertical side S of Q . This implies that $S \setminus B_k$ is disconnected. Hence there is an open subinterval I of $S \setminus B_k$ whose highest point is contained in B_k . Let D be a closed round disk centered on I and contained in $\mathbb{C} \setminus B_k$. Move the center of D upwards until the boundary of the translated disk D^* first hits B_k . Any intersection point p_k of D^* with B_k is on the top half of ∂D^* . Moreover, the unit vector \mathbf{v}_k based at p_k and pointing in the direction of the center of D^* is normal to B_k .




FIGURE 12. If the sequence of blobs has a collapsing finger, then we can find a sequence of normal vectors which accumulate to a vector which is not normal to the limiting blob.

Since we can choose Q to be arbitrarily small, we can arrange so that $p_k \rightarrow 0$. Then the normal vectors \mathbf{v}_k accumulate onto vectors pointing towards the lower half-plane at 0. This is a contradiction since the cone of normal vectors N_0 is contained in the upper half-plane. \square

Thus by Theorem 10.6, the boundary of $\text{Blob}[n, t]$ can be parametrized in a way that depends uniformly continuously on (f, t) . By Theorem 10.4, this implies that the Riemann map $G[n] : \mathbb{D} \rightarrow \text{Blob}[n, t]$ depends uniformly continuously on (f, t) . Therefore the path γ_n depends uniformly continuously on f , and hence H is continuous by Lemma 10.3. This shows that $\text{CEmb}(X, Y, h)$ is contractible whenever it is non-empty and h is generic.

11. THE REMAINING CASES

In this section, we describe the homotopy type of $\text{CEmb}(X, Y, h)$ when h is not generic. As before, we let $\{x_1, x_2, \dots\}$ be an arbitrary countable dense subset of X .

h is cyclic but not parabolic. Suppose that $h : X \rightarrow Y$ is cyclic but not parabolic. Form the annulus cover $\pi_A : \tilde{Y} \rightarrow Y$ corresponding to the image of $\pi_1(h)$. Then \tilde{Y} is isomorphic to a round annulus of finite modulus in the plane. Given $x_1 \in X$, we can define the map $\text{lift}_1 : \text{Map}(X, Y, h) \rightarrow \tilde{Y}$ and $\text{Blob}(x_1)$ in a similar way as in Section 6. The results of Sections 8 and 9 all apply in the same way except for Theorem 9.10. Let R be the maximum value of \mathfrak{m} on $\text{CEmb}(X, Y, h)$. Then $\text{Blob}(x_1, X_R)$ is homeomorphic to a circle and its complement in \tilde{Y} has two connected components. By a similar argument as in the proof of Theorem 9.10, the complement of $\text{Blob}(x_1)$ in \tilde{Y} has two connected components and hence $\text{Blob}(x_1)$ is homeomorphic to a closed annulus.

Pick any deformation retraction r of $\text{Blob}(x_1)$ into the circle $\text{Blob}(x_1, X_R)$. Given $f \in \text{CEmb}(X, Y, h)$, we let $\gamma_1(t) = \text{lift}_1(f)$ for $t \in [0, 1/2)$ and $\gamma_1(t) = r(\text{lift}_1(f), 2t-1)$ for $t \in [1/2, 1]$. Let F be the unique map in $\text{CEmb}(X_R, Y, h)$ such that $\text{lift}_1(F) = \gamma_1(1)$. The map $h[1, t] : X \setminus x_1 \rightarrow Y \setminus \pi_A(\gamma_1(t))$ is now generic so that we can construct the next paths $\gamma_2, \gamma_3, \dots$ in the same way as in Section 10. The end result is a deformation retraction of $\text{CEmb}(X, Y, h)$ into the circle $\text{CEmb}(X_R, Y, h)$.

Y is the punctured disk. Let $h : X \rightarrow \mathbb{D} \setminus 0$ be a non-trivial (hence parabolic) embedding. Since $\overline{\mathbb{D}} \setminus 0$ acts by multiplication on the space $\text{CEmb}(X, \mathbb{D} \setminus 0, h)$, the region of values $V(x_1)$ is equal to a punctured disk $r\mathbb{D} \setminus 0$ for some $r \in (0, 1)$. By Lemma 6.9, for every $y_1 \in \partial V(x_1)$ there is a unique $F \in \text{CEmb}(X, \mathbb{D} \setminus 0, h)$ such that $F(x_1) = y_1$. Given $f \in \text{CEmb}(X, \mathbb{D} \setminus 0, h)$, let γ_1 be constant equal to $f(x_1)$ on $[0, 1/2)$ followed by the radial ray from $f(x_1)$ to $rf(x_1)/|f(x_1)|$ on $[1/2, 1]$. Then construct $\gamma_2, \gamma_3, \dots$ as before (the map $h[1, t] : X \setminus x_1 \rightarrow \mathbb{D} \setminus \{0, \gamma_1(t)\}$ is generic). This gives a deformation retraction of $\text{CEmb}(X, \mathbb{D} \setminus 0, h)$ into a circle.

Y is the disk. Consider the map $\mathbb{D} \rightarrow \text{Aut}(\mathbb{D})$ which sends $a \in \mathbb{D}$ to the automorphism $M_a(z) = \frac{z-a}{1-\bar{a}z}$. Given $x_0 \in X$, we get a homeomorphism

$$\text{CEmb}(X, \mathbb{D}) \rightarrow \mathbb{D} \times \text{CEmb}(X \setminus x_0, \mathbb{D} \setminus 0)$$

given by $f \mapsto (f(x_0), M_{f(x_0)} \circ f)$. By the previous case $\text{CEmb}(X \setminus x_0, \mathbb{D} \setminus 0)$ deformation retracts into a circle. Since \mathbb{D} deformation retracts into a point, $\text{CEmb}(X, \mathbb{D})$ deformation retracts into a circle, which is homotopy equivalent to the unit tangent bundle of \mathbb{D} .

Y is the sphere. For any triple (y_1, y_2, y_3) of distinct points in $\widehat{\mathbb{C}}$, there exists a unique Möbius transformation $M[y_1, y_2, y_3]$ sending (y_1, y_2, y_3) to $(0, 1, \infty)$. Thus there is a homeomorphism

$$\text{CEmb}(X, \widehat{\mathbb{C}}) \rightarrow \text{Aut}(\widehat{\mathbb{C}}) \times \text{CEmb}(X \setminus \{x_1, x_2, x_3\}, \widehat{\mathbb{C}} \setminus \{0, 1, \infty\})$$

given by $f \mapsto (M[f(x_1), f(x_2), f(x_3)], M[f(x_1), f(x_2), f(x_3)] \circ f)$. The factor $\text{CEmb}(X \setminus \{x_1, x_2, x_3\}, \widehat{\mathbb{C}} \setminus \{0, 1, \infty\})$ deformation retracts into a point since the implicit underlying embedding is generic. As for the factor $\text{Aut}(\widehat{\mathbb{C}})$, it is homeomorphic to the set of triples (a, \mathbf{v}, b) where $a, b \in \widehat{\mathbb{C}}$ are distinct and $\mathbf{v} \in T_a \widehat{\mathbb{C}}$ is non-zero. The homeomorphism is given by $f \mapsto (f(0), f'(0), f(\infty))$. This set of triples deformation retracts into $T\widehat{\mathbb{C}}$ by moving the point b along the spherical geodesic to the antipode of a . Lastly, $T\widehat{\mathbb{C}}$ clearly deformation retracts into the unit tangent bundle $T^1\widehat{\mathbb{C}}$.

Y is the plane. Similarly, for any pair (y_1, y_2) of distinct points in \mathbb{C} , there is a unique complex affine map $M[y_1, y_2]$ sending (y_1, y_2) to $(0, 1)$. This gives a homeomorphism

$$\text{CEmb}(X, \mathbb{C}) \approx \text{Aut}(\mathbb{C}) \times \text{CEmb}(X \setminus \{x_1, x_2\}, \mathbb{C} \setminus \{0, 1\}).$$

The second factor in this product is contractible since the corresponding embedding is generic. As for the first factor, it is homeomorphic to $T\mathbb{C}$ via the map $f \mapsto (f(0), f'(0))$. The tangent bundle $T\mathbb{C}$ deformation retracts into the unit tangent bundle $T_1\mathbb{C}$, which deformation retracts further into a circle.

X is the disk. We assume that the codomain Y is not the sphere or the plane. Then we equip Y with a metric of constant curvature (either Euclidean or hyperbolic). We first define a map from the unit tangent bundle T^1Y to $\text{CEmb}(\mathbb{D}, Y)$ as follows. Given $\mathbf{v} \in T_y^1Y$, we let $D_{\mathbf{v}}$ be the largest embedded geometric disk in Y centered at y , and we let $F_{\mathbf{v}} : \mathbb{D} \rightarrow D_{\mathbf{v}}$ be the Riemann map with $F_{\mathbf{v}}(0) = y$ and $F'_{\mathbf{v}}(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. The map $\mathbf{v} \mapsto F_{\mathbf{v}}$ is an embedding. We construct a deformation retraction of $\text{CEmb}(\mathbb{D}, Y)$ into the image of that map.

Given $f \in \text{CEmb}(\mathbb{D}, Y)$, let $\mathbf{v} \in T^1Y$ be the unique vector such that $f'(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. Then let $r \in (0, 1]$ be the largest such that $f(r\mathbb{D}) \subset F_{\mathbf{v}}(\mathbb{D})$ and let $f^\dagger(z) = f(rz)$. Then $F_{\mathbf{v}}^{-1} \circ f^\dagger : \mathbb{D} \rightarrow \mathbb{D}$ is a conformal embedding with a positive derivative at the origin which it fixes.

Let $g : \mathbb{D} \rightarrow \mathbb{D}$ be a conformal embedding with $g(0) = 0$ and $g'(0) > 0$. For every $t \in (0, 1]$, define $\rho_t = \inf\{\rho > 0 : g(t\mathbb{D}) \subset \rho\mathbb{D}\}$ and $g_t(z) = g(tz)/\rho_t$. By Koebe's distortion theorem [5, p.33] we have

$$\frac{t}{(1+t)^2} \leq \frac{\rho_t}{|g'(0)|} \leq \frac{t}{(1-t)^2}$$

and it follows that $g_t \rightarrow \text{id}$ as $t \rightarrow 0$.

We define a deformation retraction of $\text{CEmb}(\mathbb{D}, Y)$ into the unit tangent bundle T^1Y by

$$H(f, t) = \begin{cases} z \mapsto f((1 - (1 - r)2t)z) & \text{if } t \in [0, 1/2) \\ F_{\mathbf{v}} \circ (F_{\mathbf{v}}^{-1} \circ f^\dagger)_{(2-2t)} & \text{if } t \in [1/2, 1]. \end{cases}$$

We leave it to the reader to check that this map is continuous.

h is trivial. Assume that Y is not the sphere or the plane, and that X is not the disk. Fix a non-zero tangent vector $\mathbf{v} \in T_x X$. Given $f \in \text{CEmb}(X, Y, h)$, we can define a disk $D_f \subset Y$ by filling the holes of $f(X)$. Then we define $F : \mathbb{D} \rightarrow D_f$ to be the Riemann map with $F'(0) = \lambda d f(\mathbf{v})$ for some $\lambda > 0$. We thus get an embedding

$$\text{CEmb}(X, Y) \rightarrow \text{CEmb}(\mathbb{D}, Y) \times \text{CEmb}(X \setminus x, \mathbb{D} \setminus 0)/S^1$$

defined by $f \mapsto (F, F^{-1} \circ f)$. The quotient $\text{CEmb}(X \setminus x, \mathbb{D} \setminus 0)/S^1$ represents those embeddings g such that $(g^{-1})'(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. There is an obvious left inverse

$$\text{CEmb}(\mathbb{D}, Y) \times \text{CEmb}(X \setminus x, \mathbb{D} \setminus 0)/S^1 \rightarrow \text{CEmb}(X, Y)$$

given by $(G, g) \mapsto G \circ g$. By previous work there is a deformation retraction H_1 of $\text{CEmb}(\mathbb{D}, Y)$ into T_1Y and a deformation retraction H_2 of $\text{CEmb}(X \setminus x, \mathbb{D} \setminus 0)/S^1$ into a point. Then $H(f, t) = H_1(F, t) \circ H_2(F^{-1} \circ f, t)$ is a deformation retraction of $\text{CEmb}(X, Y)$ into the unit tangent bundle of Y .

h is parabolic. This case is similar to the previous one. Let p be the puncture around which h wraps non-trivially and let \mathbf{v} be a non-zero vector in $T_p(Y \cup \{p\})$. Given $f \in \text{CEmb}(X, Y, h)$, we define a disk $D_f \subset Y \cup \{p\}$ by filling the holes of $f(X)$ in $Y \cup \{p\}$. Then we define $F : \mathbb{D} \rightarrow D_f$ to be the Riemann map with $F(0) = p$ and $F'(0) = \lambda \mathbf{v}$ for some $\lambda > 0$. This yields an embedding

$$\text{CEmb}(X, Y) \rightarrow (\text{CEmb}(\mathbb{D} \setminus 0, Y)/S^1) \times \text{CEmb}(X, \mathbb{D} \setminus 0)$$

defined by $f \mapsto (F, F^{-1} \circ f)$. The first factor is contractible whereas the second factor deformation retracts into a circle. We can compose the two deformation retractions to obtain a deformation retraction of $\text{CEmb}(X, Y)$ into a circle.

Y is a torus. Let Y be a torus and suppose that $h : X \rightarrow Y$ is non-trivial. Let $x_0 \in X$ and let $y_0 = h(x_0)$. For every $y \in Y$, there is a unique conformal automorphism $M_y : Y \rightarrow Y$ homotopic to the identity such that $M_y(y) = y_0$. We thus have a homeomorphism

$$\text{CEmb}(X, Y, h) \approx Y \times \text{CEmb}(X \setminus x_0, Y \setminus y_0, h).$$

given by $f \mapsto (f(x_0), M_{f(x_0)} \circ f)$. Since the restriction $h : X \setminus x_0 \rightarrow Y \setminus y_0$ is generic, the factor $\text{CEmb}(X \setminus x_0, Y \setminus y_0, h)$ is contractible. It follows that $\text{CEmb}(X, Y, h)$ is homotopy equivalent to Y .

The reader can check that we have exhausted all possibilities for the embedding $h : X \rightarrow Y$, which concludes the proof of Theorem 1.2. The latter obviously implies Theorem 1.1.

REFERENCES

- [1] L.V. Ahlfors. On quasiconformal mappings. *J. Anal. Math.*, 3(1):1–58, 1953.
- [2] L.V. Ahlfors. *Lectures on Quasiconformal Mappings: Second Edition*, volume 38. American Mathematical Society, 2006. with additional chapters by C.J. Earle and I. Kra, M. Shishikura, J.H. Hubbard.
- [3] L. Bers. *Quasiconformal mappings and Teichmüller’s theorem*. New York University, Institute of Mathematical Sciences, 1958.
- [4] M.R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer, 2009.
- [5] P.L. Duren. *Univalent functions*. Springer-Verlag, 1983.
- [6] C. Earle and A. Marden. Conformal embeddings of Riemann surfaces. *J. Anal. Math.*, 34:194–203, 1978.
- [7] C.J. Earle. Teichmüller spaces of groups of the second kind. *Acta Math.*, 112(1):91–97, 1964.
- [8] C.J. Earle. Reduced Teichmüller spaces. *Trans. Amer. Math. Soc.*, 126(1):54–63, 1967.
- [9] C.J. Earle and N. Lakic. Variability sets on Riemann surfaces and forgetful maps between Teichmüller spaces. *Ann. Acad. Sci. Fenn. Math.*, 27(2):307–324, 2002.
- [10] Y. Eliashberg and N. Mishachev. *Introduction to the h-Principle*. American Mathematical Society, 2002.
- [11] B. Farb and D. Margalit. *A Primer on Mapping Class Groups*. Princeton University Press, 2011.
- [12] J.M.G. Fell. A Hausdorff topology for the closed subsets of a locally compact non-Hausdorff space. *Proc. Amer. Math. Soc.*, 13(3):472–476, 1962.
- [13] F.P. Gardiner. Measured foliations and the minimal norm property for quadratic differentials. *Acta Math.*, 152:57–76, 1984.
- [14] F.P. Gardiner. *Teichmüller theory and quadratic differentials*. John Wiley & Sons, 1987.
- [15] F.P. Gardiner and N. Lakic. The Dirichlet principle for measured foliations. preprint, 2010.
- [16] F.P. Gardiner and N. Lakic. A synopsis of the Dirichlet principle for measured foliations. *RIMS Kôkyûroku Bessatsu*, B17:55–64, 2010.
- [17] A.A. Golubev and S.Yu. Graf. Extremal problems on classes of quasiconformal embeddings of Riemann surfaces. *Siberian Adv. Math.*, 11(4):47–64, 2001.
- [18] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [19] J.H. Hubbard. Sur les sections analytique de la courbe universelle de Teichmüller. *Mem. Amer. Math. Soc.*, 166:1–137, 1976.
- [20] M.S. Ioffe. Extremal quasiconformal embeddings of Riemann surfaces. *Sib. Math. J.*, 16:398–411, 1975.
- [21] M.S. Ioffe. Some problems in the calculus of variation “in the large” for conformal and quasiconformal imbeddings of Riemann surfaces. *Sib. Math. J.*, 19:412–424, 1978.
- [22] J.A. Jenkins. A counter-example to a statement on conformal mapping of Riemann surfaces. *Proc. Amer. Math. Soc.*, 10:423–424, 1959.
- [23] S.L. Krushkal. Variational principles in the theory of quasiconformal maps. In *Handbook of Complex Analysis: Geometric Function Theory*, volume 2, pages 31–98. Elsevier, 2005.
- [24] A. Marden and K. Strebel. The Heights theorem for quadratic differentials on Riemann surfaces. *Acta Math.*, 153:153–211, 1984.

- [25] H. Masur. On a class of geodesics in Teichmüller space. *Ann. of Math. (2)*, 102(2):205–221, 1975.
- [26] C. McMullen. Amenability, Poincaré series and quasiconformal maps. *Invent. Math.*, 97:95–127, 1989.
- [27] C. McMullen. Navigating moduli space with complex twists. *J. Eur. Math. Soc. (JEMS)*, 15:1223–1243, 2013.
- [28] J. Milnor. *Dynamics in one complex variable*. Princeton University Press, third edition, 2006.
- [29] J.R. Munkres. *Topology*. Pearson, second edition, 2000.
- [30] C. Pommerenke. *Boundary behavior of conformal maps*. Springer, 1992.
- [31] H.L. Royden. The conformal rigidity of certain subdomains on a Riemann surface. *Trans. Amer. Math. Soc.*, 76:14–25, 1954.
- [32] H.L. Royden. The interpolation problem for schlicht functions. *Ann. of Math. (2)*, 60(2):326–344, 1954.
- [33] H.L. Royden and P.M. Fitzpatrick. *Real Analysis*. Pearson, fourth edition, 2010.
- [34] M.M. Schiffer and D.C. Spencer. *Functionals of finite Riemann surfaces*. Princeton University Press, 1954.
- [35] M. Shiba. Conformal mapping of Riemann surfaces and the classical theory of univalent functions. *Publ. Inst. Math. (Beograd) (N.S.)*, 75(89):217–232, 2004.
- [36] Z. Ślodkowski. Holomorphic motions and polynomial hulls. *Proc. Amer. Math. Soc.*, 11(2):347–355, 1991.
- [37] K. Strebel. *Quadratic differentials*. Springer-Verlag, 1984.
- [38] K. Strebel. Point shift differentials and extremal quasiconformal mappings. *Ann. Acad. Sci. Fenn. Math.*, 23:475–494, 1998.
- [39] D.P. Thurston. From rubber bands to rational maps. preprint, [arXiv:1502.02561](https://arxiv.org/abs/1502.02561), 2015.
- [40] W.P. Thurston. Zippers and univalent functions. In *The Bieberbach Conjecture: Proceedings of the Symposium on the Occasion of the Proof*, pages 185–197. American Mathematical Society, 1986.

DEPARTMENT OF MATHEMATICS, THE GRADUATE CENTER, CITY UNIVERSITY OF
NEW YORK, 365 FIFTH AVENUE, NEW YORK, NY 10016, USA

E-mail address: maxforbou@gmail.com